Introduction to the Theory of Cost

Axel Liu-Bai

June 2026

Abstract

This paper serves as an expository introduction to the theory of cost. We survey the main concepts and theorems of this field. We aim to provide an intuitive view of the theory, along with accessible and detailed proofs of the results.

1 Introduction

When thinking about groups, one quantity holding information about it is the rank of the group. It gives us the minimal number of elements needed to generate the entire group. If we shift our focus to group actions, a natural question arises: is there an equivalent of the rank for group actions? In this paper, we study the notion of cost, an invariant introduced by Damien Gaboriau to study the complexity of equivalence relations and countable groups. Closely related to the study of graphs, the cost is a measure-theoretic analogue to the rank.

This paper is organized as follows. We begin by giving some motivation behind the theory of cost. Then, we build the concepts needed to define it. Finally, we present methods to simplify computations of the cost, including a proof of the Fundamental Theorem of Cost.

2 Motivation

To understand why we might be interested in the cost, consider the following example: Let $\mathbb{Z}^2 \curvearrowright \mathbb{S}^1$ be the rotation action on the circle by two generators a, b, which are rotations by rationally independent angles. This means that no matter how many times we apply a or b, we will never go back to the same point: their orbits are dense. The rank of $\mathbb{Z}^2$ is 2. However, rotations are commutative, so we do not need to define both a, b on the whole circle.

Indeed, we can define a on the whole circle and b on an ϵ-neighborhood B_ϵ of 0 (blue in Figure 1), for any $\epsilon > 0$.

Suppose we want to apply b on $x \notin B_\epsilon$. There is $n \in \mathbb{N}$ such that $a^n(x) \in B_\epsilon$, by the density of the orbit of a. Then, we can compute $a^{-n} \circ b \circ a^n(x)$, which gives us $b(x)$ without defining b on the whole circle. This is a more efficient way of describing the action than defining both generators on the whole circle.

Finally, there is a notion of orbit equivalence between group actions, which we will define later. Intuitively, it means that the orbits of the actions are the "same" but on different

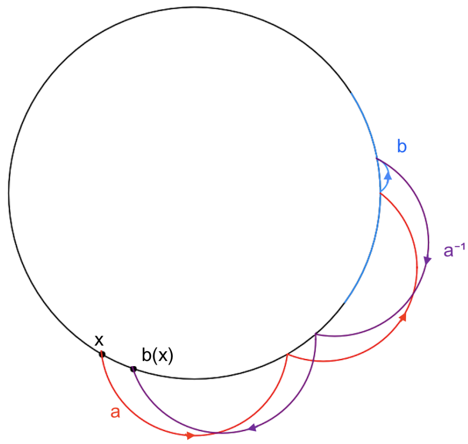

Figure 1: Computing $b(x)$ for $x \notin B_\epsilon$, with $n = 2$

spaces. Then, we would like to have a quantity that is invariant under orbit equivalence. This quantity cannot be the rank of groups: the following theorem provides counter-examples.

Theorem 1 (Ornstein, Weiss). *Any probability measure preserving action of an amenable group is orbit equivalent to a probability measure preserving action of* $\mathbb{Z}$

In particular, any probability measure preserving action of $\mathbb{Z}^d$ is orbit equivalent to a probability measure preserving action of $\mathbb{Z}$. But $\mathbb{Z}$ and $\mathbb{Z}^d$ have respective ranks 1 and d. Thus, we start the study of the cost of groups.

3 Preliminairies

In this section, we introduce the different objects we will work with. We follow chapters 1 and 2 of [Gab].

Let (X, μ) be a standard Borel atomless probability space. Let Γ be a countable group acting on (X, μ).

Definition 1 (Countable Borel equivalence relation (CBER)). *An equivalence relation* E *on* (X, μ) *is called CBER if each equivalence class is countable (there may be uncountably many classes) and* $\mathcal{R}$ *considered as a subset of* $X \times X$ *is a Borel subset.*

There are different ways of thinking about CBERs: through group actions and graphs.

3.1 Group actions

Let Γ be a countable group and α be a Borel action of Γ on (X, μ). Then, the orbit equivalence relation $\mathcal{R}_\alpha$, defined by $x\mathcal{R}_\alpha y \iff$ there is $\gamma \in \Gamma$ such that $y = \gamma \cdot x$, is a CBER. The following theorem gives us the equivalence between CBERs and Borel group actions:

Theorem 2 (Feldman-Moore). *Let* $\mathcal{R}$ *be a CBER on* X*. Then, there is a countable group* Γ *and a Borel action* α *of* Γ *on* X *such that* $\mathcal{R} = \mathcal{R}_\alpha$*. Moreover,* Γ *and* α *can be chosen so that* $x\mathcal{R}y \iff$ *there is* $\gamma \in \Gamma$ *such that* $y = \gamma \cdot x$ *and* $\gamma^2 = id$*.*

3.2 Graphs

Before diving into the equivalence between CBERs and graphs, we first need some more definitions.

Definition 2 (Probability measure preserving (pmp)). *A function $f : X \to X$ is probability measure preserving (pmp) if $f_*\mu = \mu$. $\Gamma \curvearrowright (X, \mu)$ is pmp if every $\gamma \in \Gamma$ is pmp. A CBER $\mathcal{R}$ on (X, μ) is pmp if for any Borel automorphism $f : X \to X$, i.e a Borel isomorphism with Borel inverse, such that $\forall x \in X,\ x\mathcal{R}f(x)$, we have $f_*\mu = \mu$.*

Definition 3 (Partial isomorphism). *A function $f : A \to B$ is a partial isomorphism if it is a Borel isomorphism between A and B, both Borel subsets of X.*

Definition 4 (Graphing). *A graphing is a countable collection $\Phi = (\varphi_n)_{n\in\mathbb{N}}$ of pmp partial isomorphisms of (X, μ). The equivalence relation $\mathcal{R}_\Phi$ generated by a graphing Φ is the smallest equivalence relation such that if $x \in dom(\varphi_i)$, then $xE\varphi_i(x)$.*

If $\mathcal{R}$ is a CBER on (X, μ), then a graphing Φ is a graphing of $\mathcal{R}$ if $E_\Phi = \mathcal{R}$.

Definition 5 (Schreier graph). *The Schreier graph $Sch(\mathcal{R}, \Phi)$ of $\mathcal{R}$ with respect to Φ is the directed graph defined as follows:*

- *The vertex set is $V = X$.*
- *The edge set is $E = \{(x, \varphi(x)) : \varphi \in \Phi \text{ and } x \in dom(\varphi)\}$.*

In other words, the equivalence relation $\mathcal{R}_\Phi$ is the connectedness relation of Schreier the graph $Sch(\mathcal{R}, \Phi)$. Note that $\mathcal{R}_\Phi$ is a CBER since Φ consists of Borel isomorphisms and is countable. Moreover, given a CBER $\mathcal{R}$, we get Γ and α as in theorem 2 of Feldman and Moore. Let S be a generating set for Γ. We get a graphing $\Phi = S$ since S is countable and α is Borel. This gives us the equivalence between CBERs, graphs and Borel group actions.

The Schreier graph is the equivalent of a Cayley graph, but for CBERs instead of groups. Just as different generating sets for the same group give rise to different Cayley graphs, different graphings for the same CBER give rise to different Schreier graphs. We give some examples.

Example 1 (Graphings). *Graphings that generate the equivalence relation $x \sim x + 1$ for $x \in \mathbb{R}$.*

1. *Graphing $\Phi_1 = \{\varphi_1\}$, where $\varphi_1 : \mathbb{R} \to \mathbb{R}, x \mapsto x + 1$*
2. *Graphing $\Phi_2 = \{\varphi_1, \varphi_2\}$, where $\varphi_1 : \mathbb{R} \to \mathbb{R}, x \mapsto x + 3$ and $\varphi_2 : \mathbb{R} \to \mathbb{R}, x \mapsto x - 2$*
3. *Graphing $\Phi_3 = \{\varphi_n : n \in \mathbb{Z}\}$, where $\varphi_n : [n, n + 1) \to [n + 1, n + 2), x \mapsto x + 1$*

Definition 6 (Graph of x from graphing). *Let Φ be a graphing on (X, μ) that generates a CBER $\mathcal{R}$. For each $x \in X$, we define the (rooted) directed graph $\Phi[x]$ as follows:*

- *The vertex set of $\Phi[x]$ is $V = \{y \in X : (x, y) \in \mathcal{R}\}$, the elements in the equivalence class of x.*
- *The edge set of $\Phi[x]$ is $E = \{(u, \varphi(u)) : u \in dom(\varphi), \varphi \in \Phi\}$. There is a directed edge between u, v if there is $\varphi \in \Phi$ s.t $u \in dom(\varphi)$ and $\varphi(u) = v$. Each edge $(u, \varphi(u))$ is labeled φ.*
- *The root vertex is x.*

Example 2 (Graph of x). *Recall example 1 and the graphings Φ_1, Φ_2 and Φ_3. We give an illustration of the graph of $x = 0$ from these graphings.*

1. *This graph looks the same as the Cayley graph for $\mathbb{Z}$ with generating set $S = \{1\}$.*

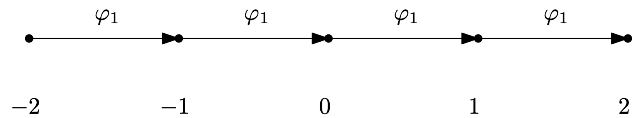

Figure 2: Graph of $x = 0$ from Φ_1

2. *This graph looks the same as the Cayley graph for $\mathbb{Z}$ with generating set $S = \{2, 3\}$.*

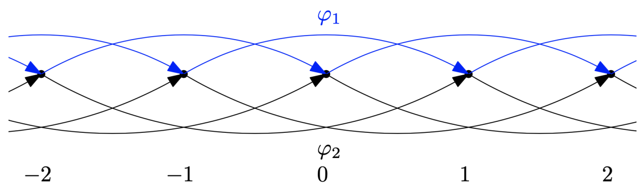

Figure 3: Graph of $x = 0$ from Φ_2

3. *This graph is similar with the one in Figure 2, but each edge comes from a different generator. The elements of Φ_3 are true partial isomorphisms, and not full isomorphisms. Each is only defined on an interval of length 1.*

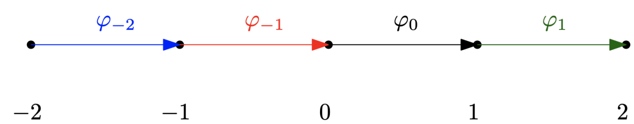

Figure 4: Graph of $x = 0$ from Φ_3

If we consider unrooted graphs, then for any $x \in X$ and $y \in \mathcal{R}(x)$, we have $\Phi[x] = \Phi[y]$. So, Φ gives a connected graph for each equivalence class. We can then define a distance d_Φ on X by using the graph distance.

Definition 7 (Metric induced by graphing). *Given a graphing Φ, we define*

$$d_\Phi(x, y) = \begin{cases} \#\ \textit{edges in the shortest path from } x \textit{ to } y \textit{ in } \Phi[x] & \textit{if } y \in \mathcal{R}(x) \\ \infty & \textit{if } y \notin \mathcal{R}(x) \end{cases}$$

Just as trees are a special case of graphs, we have the concept of treeings as a special case of graphings.

Definition 8 (Treeing). *A graphing Φ of (X, μ) is a treeing if for μ-a.e $x \in X$, the graph $\Phi[x]$ is a tree, i.e there are no cycles in the corresponding undirected graph. We say that a CBER $\mathcal{R}$ is treeable if there is a treeing Φ s.t $E_\Phi = \mathcal{R}$.*

We also have the notion of total degree of a vertex in graphings, which is the valency.

Definition 9 (Valency). *The valency of $x \in X$ in $\Phi[x]$ is the number of domains and images $dom(\varphi_n), im(\varphi_n)$ that x is contained it. More specifically,*

$$v_\Phi(x) = \sum_{n\in\mathbb{N}} \left(\mathbb{1}_{dom(\varphi_n)} + \mathbb{1}_{dom(\varphi_n)}\right)(x)$$

4 Cost

In this section, we introduce the cost, following chapter 2 of [Gab]. We have different definitions for the different ways of representing a CBER, but the definitions build upon each other.

Definition 10 (Cost of a graphing Φ). *The cost $\mathcal{C}(\Phi)$ of Φ is the number of generators φ_n weighted by the measure of their support, which is their domain $dom(\varphi_n)$:*

$$\mathcal{C}(\Phi) = \sum_{n\in\mathbb{N}} \mu(dom(\varphi_n))$$

Since φ_n is an isomorphism between $dom(\varphi_n)$ and $im(\varphi_n)$, and it is pmp, then $\mu(dom(\varphi_n)) = \mu(im(\varphi_n))$, so we also have

$$\mathcal{C}(\Phi) = \sum_{n\in\mathbb{N}} \mu(im(\varphi_n))$$

We can also write the cost in terms of valency: By integrating v_Φ, we obtain

$$\begin{aligned}
\int_X v_\Phi \, d\mu &= \int_X \left(\sum_{n\in\mathbb{N}} \mathbb{1}_{dom(\varphi_n)} + \mathbb{1}_{im(\varphi_n)}\right) \, d\mu \\
&= \sum_{n\in\mathbb{N}} \int_X \mathbb{1}_{dom(\varphi_n)} + \mathbb{1}_{im(\varphi_n)} \, d\mu && \text{ctbl add. of } \int \\
&= \sum_{n\in\mathbb{N}} \mu(dom(\varphi_n)) + \mu(im(\varphi_n)) \\
&= 2\mathcal{C}(\Phi)
\end{aligned}$$

Hence,

$$\mathcal{C}(\Phi) = \frac{1}{2}\int_X v_\Phi \, d\mu$$

We this definition, we can think of the cost of graphs as half the average degree of vertices.

Definition 11 (Cost of an equivalence relation $\mathcal{R}$). *The cost of an equivalence relation $\mathcal{R}$ is the infimum of the cost of graphings generating $\mathcal{R}$:*

$$\mathcal{C}(\mathcal{R}) = \inf\{\mathcal{C}(\Phi) : E_\Phi = \mathcal{R}\} = \inf\{\nu(A) : A \subseteq \mathcal{R}, A \text{ generates } \mathcal{R}\}$$

This is an OE-invariant, i.e if $\mathcal{R}_1$ is OE to $\mathcal{R}_2$, then $\mathcal{C}(\mathcal{R}_1) = \mathcal{C}(\mathcal{R}_2)$. Indeed, let $f : X_1 \to X_2$ be the Borel bijection in the definition of OE. For Φ_1 s.t $E_{\Phi_1} = \mathcal{R}_1$, we construct Φ_2 s.t $E_{\Phi_2} = \mathcal{R}_2$ and $\mathcal{C}(\Phi_1) = \mathcal{C}(\Phi_2)$. Let $\varphi_1 \in \Phi_1$. Then, $\varphi_2 = f \circ \varphi_1 \circ f^{-1} : f(dom(\varphi_1)) \to$

$f(im(\varphi_1))$ is a partial isomorphism since f is a Borel bijection. Let $\Phi_2 = \{\varphi_2 = f \circ \varphi_1 \circ f^{-1} : \varphi_1 \in \Phi_1\}$. This generates the equivalence relation $\mathcal{R}_2$ since f preserves equivalence classes: If $f(x) \sim_2 f(y)$, then $x \sim_1 y$ so there is φ_1 s.t $\varphi_1(x) = y$. Then, we have

$$\varphi_2(f(x)) = f \circ \varphi_1 \circ f^{-1}(f(x)) = f(\varphi(x)) = f(y)$$

The costs are equal since f is pmp: For each $\varphi_1 \in \Phi_1$, we have

$$\mu_2(dom(\varphi_2)) = f_*\mu_1(f(dom(\varphi_1))) = \mu_1(dom(\varphi_1))$$

Hence, $\mathcal{C}(R_1) = \mathcal{C}(R_2)$.

Finally, we define the cost of a group in terms of the cost of the equivalence relations generated by their free pmp actions.

Definition 12 (Cost of a group Γ). *The cost $\mathcal{C}_*(\Gamma)$ of a group is the infimum of the cost of its free pmp actions:*

$$\mathcal{C}_*(\Gamma) = \inf\{\mathcal{C}(\mathcal{R}_\alpha) : \alpha \textit{ is free and pmp}\}$$

It is also called the infimum cost or minimal cost. The sup-cost $\mathcal{C}^*(\Gamma)$ of a group is the supremum of the cost of its free pmp actions:

$$\mathcal{C}_*(\Gamma) = \sup\{\mathcal{C}(\mathcal{R}_\alpha) : \alpha \textit{ is free and pmp}\}$$

A group Γ has fixed price if $\mathcal{C}_*(\Gamma) = \mathcal{C}^*(\Gamma)$.

We will focus on the infimum cost in the rest of this paper. We observe that the computation of the cost for equivalence relations and groups is quite difficult because of the definition with infimums. We will now see two methods that help simplify the computation: the induction formula and the Fundamental Theorem of Cost.

5 Induction Formula

The induction formula will allow us to compute the cost of the entire equivalence relation $\mathcal{R}$ on X by only computing the cost on a smaller appropriate subset Y of X. We follow section 2 of [Gab00] and give an English translation of the proof.

5.1 Preliminairies

Definition 13 (Induced equivalence relation). *Let $\mathcal{R}$ be a CBER on X and $Y \subseteq X$ be a Borel set of positive measure. Then, we get an equivalence relation $\mathcal{R}_{|Y}$ on (Y, μ_Y) induced by $\mathcal{R}$: For $x, y \in Y$, $x\mathcal{R}_{|Y}y \iff x\mathcal{R}y$.*

Definition 14 (Fundamental domain). *A fundamental domain for an equivalence relation on X is a Borel subset $D \subseteq X$ that meets almost each class exactly once (up to a countable union of null equivalence classes).*

A fundamental domain is a transversal of the equivalence relation, up to a null set.

Definition 15 (Complete section). *A complete section for an equivalence relation on X is a Borel subset $C \subseteq X$ that meets almost each class.*

Definition 16 (Φ-word). *Let Φ be a graphing on X. A Φ-word ω is a sequence of elements of Φ or their inverse: $\omega = \varphi_{i_l}^{\epsilon_l} \dots \varphi_{i_1}^{\epsilon_1}$, where $\varphi_{i_j} \in \Phi$ and $\epsilon_j = \pm 1$. A word is reduced if there are no $j \in \{1, \dots, l-1\}$ s.t $i_j = i_{j+1}$ and $\varphi_{i_j}^{-1} = \varphi_{j+1}$. Then, the length of w is l.*

If $Y \subseteq X$ meets each equivalence class, then for $x \in X$, a Φ-word ω of x at Y is a word s.t $\omega(x) \in Y$.

5.2 Induction Formula

Proposition 1 (Induction). *Let $\mathcal{R}$ be a CBER on (X, μ) and $Y \subseteq X$ be a Borel subset intersecting each class of $\mathcal{R}$. Then,*

1. *$\mathcal{R}$ is treeable $\iff$ $\mathcal{R}_{|Y}$ is treeable*
2. *If $0 < \mu(X) < \infty$, let $\mu_1 = \frac{\mu}{\mu(X)}$ and $\nu_1 = \frac{\mu_{|Y}}{\mu(Y)}$ be probability measures on X and Y respectively, then*

$$\mu(X) \cdot (\mathcal{C}_{\mu_1}(\mathcal{R}) - 1) = \mu(Y) \cdot (\mathcal{C}_{\nu_1}(\mathcal{R}_{|Y}) - 1)$$

The proof of this proposition requires a lemma.

Lemma 1. *Let Φ be a graphing of the CBER $\mathcal{R}$ on (X, μ), $Y \subseteq X$ be a Borel complete section and $\mathcal{R}_{|Y}$ be the induced equivalence relation from $\mathcal{R}$. Then, by splitting the domains of the generators if needed, we can partition Φ in two disjoint subgraphings $\Phi = \Phi_v \sqcup \Phi_h$ (vertical, horizontal) s.t*

1. *Φ_v is a treeing where each orbit intersects Y in exactly one point.*
2. *"Sliding" Φ_h along Φ_v into Y gives a graphing Ψ_h of $\mathcal{R}_{|Y}$. ("Sliding" is defined in the proof)*
3. (a) *If Ψ' is another graphing of $\mathcal{R}_{|Y}$, then $\Phi' = \Phi_v \sqcup \Psi'$ is a graphing of $\mathcal{R}$.*
   (b) *If Φ is a treeing of $\mathcal{R}$, then Ψ_h is a treeing of $\mathcal{R}_{|Y}$. If Ψ' is a treeing of $\mathcal{R}_{|Y}$, then Φ' is a treeing of $\mathcal{R}$.*

If $\mu(X) < \infty$, then the orbits of Φ_v are finite and

4. *$\mathcal{C}_\mu(\Phi_v) = \mu(X) - \mu(Y)$*
5. *$\mathcal{C}_\mu(\Phi_h) = \mathcal{C}_{\mu_{|Y}}(\Psi_h)$*

Proof. We prove this lemma in the following order: 1,4,5,2,3.

Recall that a graphing Φ gives a distance d_Φ on X defined by shortest paths. Let $Y_0 = Y$ and for $i \in \mathbb{N}$, let $Y_i = \{x \in X : d_\Phi(x, Y_0) = i\}$, where $d_\Phi(x, Y_0) := \min_{y \in Y_0} d_\Phi(x, y)$.

By taking unions of open balls $B_{i+1}(y)$ for $y \in Y$ and then intersecting with Y_{i-1}^c, we obtain by induction that the Y_i are Borel since $Y_0 = Y$ is Borel. For each $\varphi \in \Phi$ and $i \in \mathbb{N}$, we have that $dom(\varphi) \cap Y_i$ and $im(\varphi) \cap Y_i$ are Borel, countable and generate the same equivalence relation. Moreover, by symmetry, φ and φ^{-1} relate the same elements.

Thus, by splitting the domains and images of the generators and replacing φ with φ^{-1}, we can assume wlog that for each $\varphi : A \to B$ in Φ, we have $A \subseteq Y_i,\ B \subseteq Y_j,\ i \geq j \in \mathbb{N}$. In fact $j = i - 1$ or $j = 1$ since for $x \in A$, there is an edge between x and $\varphi(x)$. So $d_\Phi(x, Y) \leq d_\Phi(\varphi(x)) + 1$ and $d_\Phi(\varphi(x), Y) \leq d_\Phi(x) + 1$

Visualizing Y_0 at the "bottom" of X, we have that each φ maps towards the bottom, or equivalently, the domain is "higher" than the image.

We number the φ (starting at 1), then for $i \in \mathbb{N}^+$ and $\varphi_j \in \Phi$, let

$$Z_i^{\varphi_j} = \{x \in Y_i : \varphi_j \text{ is the first generator to map } x \text{ into } Y_{i-1}\}$$

Since Y intersects every equivalence class and $d_\Phi(x) = i$, there is a Φ-word ω of x at Y that has length i. Each letter of ω has domain in Y_j and image in Y_{j-1}. So the first generator in ω maps x into Y_{i-1}. Thus, such a generator exists and we can pick the first one to be in $Z_i^{\varphi_j}$.

This gives us a partition of $X \setminus Y$:

$$X \setminus Y = \bigsqcup_{i,j\in\mathbb{N}^+} Z_i^{\varphi_j} \tag{$*$}$$

We now define Φ_v, Φ_h: For $\varphi : A \subseteq Y_i \to B$, let $\delta'_\varphi = \varphi_{|Z_i^\varphi}$ and $\delta_\varphi = \varphi_{|dom(\varphi)\setminus Z_i^\varphi}$. Intuition: This means if $B \subseteq Y_i$ (horizontal edges), then $\delta'_\varphi = 0$ and $\delta_\varphi = \varphi$. If $B \subseteq Y_{i-1}$ (vertical or diagonal edges), then δ'_φ gets the first vertical edges and δ_φ gets all the other diagonal edges. Then,

$$\Phi_v = \{\delta'_\varphi : \varphi \in \Phi\} \qquad\qquad \Phi_h = \{\delta_\varphi : \varphi \in \Phi\}$$

Since the Z_i^φ partition $X \setminus Y$, then for each $x \in X \setminus Y$, there is a unique generator δ'_{φ_j} in Φ_v s.t $x \in dom(\delta'_{\varphi_j})$. Moreover, φ_j maps down towards Y_{j-1}, i.e $im(\delta'_{\varphi_j}) \subseteq Y_{j-1}$. Thus, there is a unique reduced Φ_v-word ω_x of x at Y. This uniqueness gives us that each orbit of Φ_v is a tree that meets Y in exactly one point, which is the root. This proves part 1.

By part 1, Y is a fundamental domain for E_{Φ_v}. Then, for a group element γ, we have $(\gamma \cdot Y) \cap Y = \emptyset$. Note that since $\mathcal{R}$ is a CBER, then so is E_{Φ_v}. So $\mu(\gamma \cdot Y) = \mu(Y)$. Then,

$$|orb(\Phi_v)| \cdot \mu(Y) = \sum_{\gamma\in orb(\Phi_v)} \mu(\gamma \cdot Y) = \mu\left(\bigsqcup_{\gamma\in orb(\Phi_v)} \mu(\gamma \cdot Y)\right) \leq \mu(X)$$

If $\mu(X) < \infty$, then the orbits of Φ_v must be finite. Moreover, by $(*)$ and finiteness, we have

$$\mu(X) - \mu(Y) = \mu(X \setminus Y) = \sum_{i,j\in\mathbb{N}^+} \mu(Z_i^{\varphi_j}) = \sum_{\delta'\in\Phi_v} \mu(dom(\delta')) = \sum_{j\in\mathbb{N}} \mu(dom(\varphi_j))$$

This proves part 4.

By our argument above for part 1, we have that for each $\delta \in \Phi_h$, $\delta : A \to B$, and each $x \in A$ (resp. $y \in B$), there is a unique Φ_v word ω_x of x (resp. ω_y of y) at Y. By splitting the domain and image of δ, we can assume that for each $\delta \in \Phi_h$, $\delta : A \to B$, there is a unique Φ_v word $\omega_{A,\delta}$ (resp. $\omega_{B,\delta}$) s.t $\forall x \in A$ (resp. $y \in B$), it is the Φ_v-word of x (resp. y) at Y. Note that we can do the splitting by intersecting the domain (resp. image) of δ, which is Borel, with the preimage of Y wrt. ω_x (resp. ω_y).

Moreover, if A (resp. B) is contained in Y, then the word is Id. Now, for $\delta \in \Phi_h$, let

$$\psi_\delta = \omega_{B,\delta} \circ \delta \circ (\omega_{A,\delta})^{-1} : \omega_{A,\delta}(A) \to \omega_{B,\delta}(B)$$

and let Ψ_h be the collection of the ψ_δ. "Sliding" Φ_h along Φ_v means we map $\delta \in \Phi_h$ to $\psi_\delta \in \Psi_h$. Since $\omega_{A,\delta}(A) \subseteq Y$, then $\mu_{|Y}(\omega_{A,\delta}(A)) = \mu(\omega_{A,\delta}(A))$. The word $\omega_{A,\delta}$ is a composition of partial isomorphisms so $\mu(\omega_{A,\delta}(A)) = \mu(A)$. Then, for each $\delta \in \Phi_h$, we get

$$\mu_{|Y}(dom(\psi_\delta)) = \mu_{|Y}(\omega_{A,\delta}(A)) = \mu(\omega_{A,\delta}(A)) = \mu(A) = \mu(dom(\delta))$$

Thus, we have proved part 5:

$$\mathcal{C}_\mu(\Phi_h) = \sum_{\delta\in\Phi_h} \mu(dom(\delta)) = \sum_{\psi_\delta\in\Psi_h} \mu_{|Y}(dom(\psi_\delta)) = \mathcal{C}_{\mu_{|Y}}(\Psi_h)$$

By our construction/assumptions, we represent the graphing Φ as a forest of trees with edges pointing downward, root in Y and additional horizontal edges between trees. Then, we can think of Φ_v as all the vertical edges and Φ_h as all the horizontal edges. Ψ_h is sliding a horizontal edge down its endpoints' trees to reach Y. This preserves the equivalence relation on Y, so Ψ_h is a graphing of $\mathcal{R}_{|Y}$, which proves part 2.

If Ψ' is another graphing of $\mathcal{R}_Y$, then $\Phi' = \Phi_v \sqcup \Psi'$ adds back the trees, so the rest of the equivalence relation. Hence Φ' is a graphing of $\mathcal{R}$, which proves part 3.a.

If Φ is a treeing of $\mathcal{R}$, then there is at most one horizontal edge between two trees and no horizontal edge within a single tree. Then, sliding this edge down to Y does not create a cycle so Ψ_h is a treeing of $\mathcal{R}_{|Y}$. If Ψ_h is a treeing of $\mathcal{R}_{|Y}$, then adding back the vertical trees does not add cycles so Φ' is a treeing of $\mathcal{R}$. This proves 3.b. $\square$

Proof (Proposition 1). We prove this using Lemma 1:

1. This follows immediately from part 3.b of the lemma
2. First, note that $\mathcal{C}_\mu(\cdot) = \mu(X)\mathcal{C}_{\mu_1}(\cdot)$ and similarly, $\mathcal{C}_{\mu_{|Y}}(\cdot) = \mu(Y)\mathcal{C}_{\nu_1}(\cdot)$. Now, since $0 < \mu(X) < \infty$, then we have

$$\mathcal{C}_{\mu_{|Y}}(\Psi_h) = \mathcal{C}_\mu(\Phi_h) = \mathcal{C}_\mu(\Phi) - \mathcal{C}_\mu(\Phi_v) = \mu(X)\cdot\mathcal{C}_{\mu_1}(\Phi) - (\mu(X) - \mu(Y))$$

where we use part 5, $\Phi = \Phi_v \sqcup \Phi_h$ and part 4 of the lemma. Since Ψ_h is a graphing of $\mathcal{R}_{|Y}$, then

$$\mu(Y)\cdot\mathcal{C}_{\nu_1}(\mathcal{R}_{|Y}) \leq \mu(Y)\cdot\mathcal{C}_{\nu_1}(\Psi_h) = \mathcal{C}_{\mu_{|Y}}(\Psi_h)$$

Thus, combining these, we get

$$\mu(Y)\cdot\mathcal{C}_{\nu_1}(\mathcal{R}_{|Y}) \leq \mu(X)\cdot\mathcal{C}_{\mu_1}(\Phi) - (\mu(X) - \mu(Y))$$

So taking infimums over graphings Φ of $\mathcal{R}$, we get

$$\mu(Y)\cdot\mathcal{C}_{\nu_1}(\mathcal{R}_{|Y}) \leq \mu(X)\cdot\mathcal{C}_{\mu_1}(\mathcal{R}) - (\mu(X) - \mu(Y)) \tag{$*$}$$

Now, if Ψ' is a graphing of $\mathcal{R}_{|Y}$, then by part 3.a of the lemma, $\Phi' = \Phi_v \sqcup \Psi'$ is a graphing of $\mathcal{R}$. Then,

$$\mu(X)\cdot\mathcal{C}_{\mu_1}(\mathcal{R}) \leq \mu(X)\cdot\mathcal{C}_{\mu_1}(\Phi') = \mu(X)\cdot\mathcal{C}_{\mu_1}(\Phi_v) + \mu(X)\cdot\mathcal{C}_{\mu_1}(\Psi')$$

Using part 4 of the lemma with $\mathcal{C}_{\mu_1}(\Phi_v)$, we get

$$\begin{aligned}
\mu(X)\cdot\mathcal{C}_{\mu_1}(\mathcal{R}) &\leq \mu(X)\cdot(\mu_1(X)-\mu_1(Y))+\mu(X)\cdot\mathcal{C}_{\mu_1}(\Psi') && \text{by 4}\\
&= (\mu(X)-\mu(Y))+\mathcal{C}_{\mu}(\Psi')\\
&= (\mu(X)-\mu(Y))+\mathcal{C}_{\mu_{|Y}}(\Psi') && \Psi' \text{ graphing of } \mathcal{R}_{|Y}\\
&= (\mu(X)-\mu(Y))+\mu(Y)\cdot\mathcal{C}_{\nu_1}(\Psi')
\end{aligned}$$

Rearranging, we obtain

$$\mu(X)\cdot\mathcal{C}_{\mu_1}(\mathcal{R})-(\mu(X)-\mu(Y))\leq \mu(Y)\cdot\mathcal{C}_{\nu_1}(\Psi')$$

So taking infimums over graphings Ψ' of $\mathcal{R}_{|Y}$, we get

$$\mu(X)\cdot\mathcal{C}_{\mu_1}(\mathcal{R})-(\mu(X)-\mu(Y))\leq \mu(Y)\cdot\mathcal{C}_{\nu_1}(\mathcal{R}_{|Y}) \tag{$**$}$$

We have proved both inequalities, $(*)$ and $(**)$. Thus,

$$\begin{aligned}
&\mu(X)\cdot\mathcal{C}_{\mu_1}(\mathcal{R})-(\mu(X)-\mu(Y))=\mu(Y)\cdot\mathcal{C}_{\nu_1}(\mathcal{R}_{|Y})\\
\implies\quad &\mu(X)\cdot\mathcal{C}_{\mu_1}(\mathcal{R})-\mu(X)=\mu(Y)\cdot\mathcal{C}_{\nu_1}(\mathcal{R}_{|Y})-\mu(Y)\\
\implies\quad &\mu(X)\cdot(\mathcal{C}_{\mu_1}(\mathcal{R})-1)=\mu(Y)\cdot(\mathcal{C}_{\nu_1}(\mathcal{R}_{|Y})-1)
\end{aligned}$$

□

6 Fundamental Theorem of Cost

We present and prove the Fundamental Theorem of Cost, following the understanding of Jenna Zomback found in [Zom20].

Recall that the cost for equivalence relations and groups is difficult to compute through their definition, but the one for graphings is much more hands on. Thinking about graphs will give us a good idea on the optimal way to compute the cost of $\mathcal{R}_\alpha$ for a group action α. All the information we are concerned about is contained in the connectedness relation.

How would we describe the connectedness in a graph using the least amount of information possible? A spanning tree is the most efficient way by avoiding any redundancies, which are the cycles. Coming back to CBERs, we would expect that a treeing would useful for computing the cost. The Fundamental Theorem of Cost confirms our intuition.

Theorem 3 (FTC). *If Ψ is a treeing of a pmp CBER $\mathcal{R}$, then $\mathcal{C}(\Psi)=\mathcal{C}(\mathcal{R})$*

6.1 Reductions

We will reduce the proof first. Let Ψ be a treeing of $\mathcal{R}$ and Φ be a graphing of $\mathcal{R}$. Let T be the Schreier tree of Ψ and G be the Schreier graph of Φ. Let $\mathcal{C}(T)=\mathcal{C}(\Psi)$ and $\mathcal{C}(G)=\mathcal{C}(\Phi)$. Since the cost of $\mathcal{R}$ is defined as an infimum over graphings, it is enough to show that $\mathcal{C}(T)\leq\mathcal{C}(G)$. If $\mathcal{C}(G)=\infty$, we are done so suppose $\mathcal{C}(G)<\infty$.

Reduction 1. *We can assume $\mathcal{C}(T)<\infty$. In fact, we can assume T has bounded degree*

Proof. We prove that if FTC holds when T has bounded degree, then we can prove it for any T. By the Feldman-Moore theorem, there is a countable group Γ and a Borel action of Γ on X s.t $E_T = E_\Gamma$ and $xTy \iff \exists \gamma \in \Gamma$ s.t $\gamma \cdot x = y$. Let $S = \{\gamma_n\}$ be a set of generators of Γ. We can get a Borel coloring $c : T \to \mathbb{N}$ of the edges of T by

$$c(e) = n \in \mathbb{N} \text{ s.t } \gamma_n \cdot x = y, \text{ where } e = (x, y)$$

Let $T_n = \{e \in T : c(e) \leq n\}$, so $T = \bigcup_{n\in\mathbb{N}} T_n$ is an increasing union and it is enough to show $\mathcal{C}(T_m) \leq \mathcal{C}(G)$ for each $m \in \mathbb{N}$. Note that each T_n has finite cost since we only include a finite number of generators (finite groups have finite cost). In fact, T_n has bounded degree.

Let $n > m$. Let $G_n = G \cap E_{T_n}$, so $G = \bigcup_{n\in\mathbb{N}} G_n$ is an increasing union. Let $G'_n = G_n \sqcup (T_n \setminus E_{G_n})$, which is a graphing of E_{T_n}. Indeed, for xT_ny, the path between x, y is either in G_n or in $T_n \setminus E_{G_n}$. So T_n is a treeing of E_{T_n}, G'_n is a graphing of E_{T_n} and T_n has bounded degree. By FTC with bounded degree, we have $\mathcal{C}(T_n) \leq \mathcal{C}(G'_n)$. Thus, since $m < n$, we have

$$\mathcal{C}(T_n) = \mathcal{C}(T_n \setminus T_m) + \mathcal{C}(T_m) \tag{1}$$

and (2):

$$\begin{aligned}
\mathcal{C}(G'_n) &= \mathcal{C}(G_n) + \mathcal{C}(T_n \setminus E_{G_n}) && \text{def.}\\
&= \mathcal{C}(G_n) + \mathcal{C}(T_n \setminus (T_m \cup E_{G_n})) + \mathcal{C}(T_m \setminus E_{G_n}) && \text{add.}\\
&\leq \mathcal{C}(G_n) + \mathcal{C}(T_n \setminus T_m) + \mathcal{C}(T_m \setminus E_{G_n}) && \text{monotonicity}
\end{aligned} \tag{2}$$

Combining these, we get

$$\begin{aligned}
\mathcal{C}(T_m) &= \mathcal{C}(T_n) - \mathcal{C}(T_n \setminus T_m) && (1)\\
&\leq \mathcal{C}(G'_n) - \mathcal{C}(T_n \setminus T_m) && \text{FTC}\\
&\leq \mathcal{C}(G'_n) + \mathcal{C}(G_n) + \mathcal{C}(T_m \setminus E_{G_n}) - \mathcal{C}(G'_n) && (2)\\
&= \mathcal{C}(G_n) + \mathcal{C}(T_m \setminus E_{G_n})
\end{aligned}$$

Since $n > m$ was arbitrary, taking $n \to \infty$, we obtain $E_{G_n} \to E$ and $G_n \to G$. Then, $T_m \setminus E_{G_n} \to T_m \setminus E = \emptyset$. Thus, $\mathcal{C}(G_n) + \mathcal{C}(T_m \setminus E_{G_n}) \to \mathcal{C}(G)$ and so $\mathcal{C}(T_m) \leq \mathcal{C}(G)$. □

Recall Definition 7 of the metric induced by graphings.

Reduction 2. *We can assume there are $L, M > 0$ s.t*

$$\frac{1}{M} d_T \leq d_G \leq L d_T \tag{$*$}$$

Proof. We prove that if FTC holds when there are such $L, M > 0$ satisfying $(*)$, then we can prove it for any T. By Reduction 1, we can assume G, T have finite cost. It is enough to show that $\forall \epsilon > 0$, there is a graphing G'' of E and $L, M > 0$ s.t $\mathcal{C}(G \triangle G'') < \epsilon$ and $(*)$ holds. Indeed, we would get $\mathcal{C}(T) \leq \mathcal{C}(G'')$ so $\mathcal{C}(T) \leq_\epsilon \mathcal{C}(G)$ for all $\epsilon > 0$. Hence, $\mathcal{C}(T) \leq \mathcal{C}(G)$.

Let $\epsilon > 0$ and define $G_n := \{(x, y) \in G : d_T(x, y) \leq n\}$ and $T_n := T \cap E_{G_n}$. Then, G, T are increasing unions: $G = \bigcup_{n\in\mathbb{N}} G_n$ and $T = \bigcup_{n\in\mathbb{N}} T_n$. Pick n large enough so that $\mathcal{C}(T \setminus T_n) < \epsilon/4$ and $\mathcal{C}(G \setminus G_n) < \epsilon/4$. Now, let

$$G' := G_n \sqcup (T \setminus T_n) = G_n \sqcup (T \setminus E_{G_n})$$

With a similar reasoning as in Reduction 1, G' is a graphing of E. Moreover, we have

$$G \setminus G' = G \setminus (G_n \sqcup (T \setminus E_{G_n})) = G \setminus G_n \cap G \setminus (T \setminus E_{G_n}) \subseteq G \setminus G_n$$

and

$$G' \setminus G = (G_n \sqcup (T \setminus T_n)) \setminus G = (G_n \setminus G) \sqcup ((T \setminus T_n) \setminus G) = (T \setminus T_n) \setminus G \subseteq T \setminus T_n$$

Thus,

$$\mathcal{C}(G \triangle G') = \mathcal{C}(G \setminus G') + \mathcal{C}(G' \setminus G) \leq \mathcal{C}(G \setminus G_n) + \mathcal{C}(T \setminus T_n) < \epsilon/2 \qquad (*)$$

We claim that $d_T \leq n d_{G'}$: Suppose $xG'y$ so $d_{G'}(x, y) = 1$. We show that $d_T(x, y) \leq n$. If xG_ny, then by definition of G_n, we have $d_T(x, y) \leq n$. If $x(T \setminus T_n)y$, then in particular xTy, so $d_T(x, y) = 1 \leq n$. Thus, set $M := n$.

Now, for the other inequality. Let $T_n^m := \{(x, y) \in T_n : d_{G_n}(x, y) \leq m\}$ so $T_n = \bigcup_{m \in \mathbb{N}} T_n^m$ is an increasing union. Take m large enough so that $\mathcal{C}(T_n \setminus T_n^m) < \epsilon/2$. Let

$$G'' := G' \cup (T_n \setminus T_n^m)$$

Since G' was a graphing of E, so is G''. Moreover, we have

$$G \setminus G'' = G \setminus (G' \cup (T_n \setminus T_n^m)) = (G \setminus G') \cap (G \setminus (T_n \setminus T_n^m)) \subseteq G \setminus G'$$

and

$$G'' \setminus G = (G' \cup (T_n \setminus T_n^m)) \setminus G = (G' \setminus G) \cup ((T_n \setminus T_n^m) \setminus G) \subseteq (G' \setminus G) \cup (T_n \setminus T_n^m)$$

Thus,

$$\begin{aligned}
\mathcal{C}(G \triangle G'') &= \mathcal{C}(G \setminus G'') + \mathcal{C}(G'' \setminus G) && \text{def.} \\
&\leq \mathcal{C}(G \setminus G') + \mathcal{C}((G' \setminus G) \cup (T_n \setminus T_n^m)) && \text{monotonicity} \\
&\leq \mathcal{C}(G \setminus G') + \mathcal{C}(G' \setminus G) + \mathcal{C}(T_n \setminus T_n^m) && \text{subadd.} \\
&= \mathcal{C}(G \triangle G') + \mathcal{C}(T_n \setminus T_n^m) && \text{add.} \\
&< \epsilon/2 + \epsilon/2 && (*) \text{ and choice of } m \\
&= \epsilon
\end{aligned}$$

We still have $d_T \leq n d_{G''}$ since $d_T \leq n d_{G'}$ and $T_n \setminus T_n^m \subseteq T$. We claim that $d_{G''} \leq m d_T$: Suppose xTy so $d_T(x, y) = 1$. We show that $d_{G''}(x, y) \leq m$. If xT_n^my, then by definition of T_n^m, we have $d_G(x, y) \leq m$. If $x(T \setminus T_n)y$, then in particular $xG'y$ so $xG''y$ and $d_{G''}(x, y) = 1 \leq m$. Finally, if $x(T_n \setminus T_n^m)y$, then in particular $xG''y$ so $d_{G''}(x, y) = 1 \leq m$. We have treated all the possibilities of xTy. Thus, set $L := m$. G'' is as desired. □

6.2 Proof of FTC

We are now ready to prove FTC. We will illustrate the constructions in the proof by the following example on the treeing T and graphing G:

(a) Treeing T (b) Graphing G

Figure 5: Example

Proof of FTC. By Reductions 1 and 2, we can assume wlog that T has bounded degree and there are $L, M > 0$ s.t $\frac{1}{M} d_T \leq d_G \leq L d_T$. Let $d < \infty$ bound the degree of T. We define the Euler characteristic of a graphing G on (X, μ) by $Euler_\mu(G) = \mathcal{C}(G) - \mu(X)$. To show $\mathcal{C}(T) \leq \mathcal{C}(G)$, it is equivalent to showing $Euler_\mu(T) \leq Euler_\mu(G)$ since G and T are on the same space.

To relate the graphing G to the treeing T, we will first implement G with the edges of T using the fact that between two vertices in a tree, there is a unique path between them. Call this implemented graph G', which is built as follows: The vertex set will be X'.

1. For each vertex $x \in X$, create a vertex $(x, \emptyset) \in X'$
2. For each edge $e = (u, v) \in G$,
    - If $e \in T$, then add the edge $[(u, \emptyset), (v, \emptyset)]$ to G'
    - If $e \notin T$, let $u = x_0, x_1, \ldots, x_n = v$ be the path in T between u, v. Create fake vertices $(x_1, e), \ldots, (x_{n-1}, e) \in X'$. Add the path $(x_0, \emptyset), (x_1, e), \ldots, (x_{n-1}, e), (x_n, \emptyset)$ to G'.

Formally,

$$X' = (X \times \{\emptyset\}) \sqcup \{(x, e) : x \text{ lies in the interior of the path in } T \text{ between the ends of } e\}$$

Example 3. *To implement the edges of G with the edges of T, we will need to do some extra work for the edges e, f in red and blue respectively in Figure 6.*

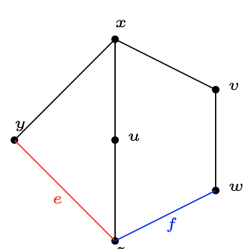

Figure 6: Edges to implement

The black edges are those already in T. We implement those first:

The red edge e has endpoints y, z and the path from y to z in T goes through x and u. We create copies (x, e) and (u, e) for these vertices and we get

Finally, the blue edge f has endpoints w, z and the path from w to z in T goes through v, x and u. We create copies of these vertices with second coordinate f and we get

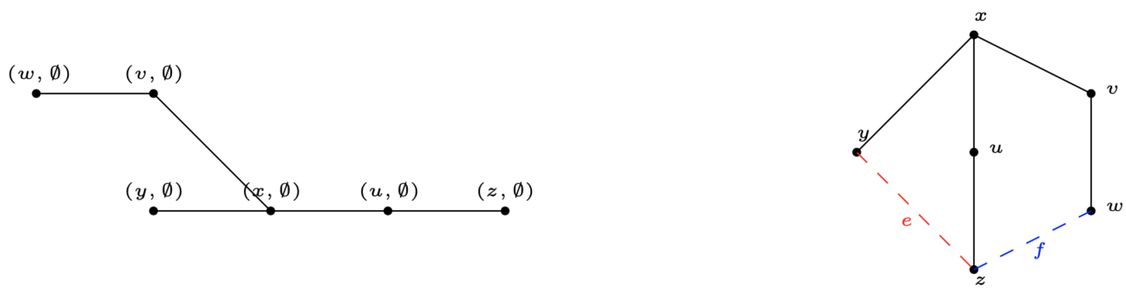

Figure 7: Step 1 of building G'

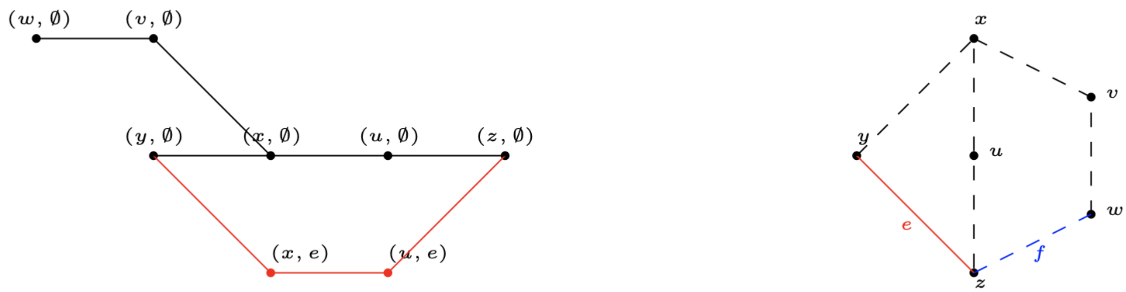

Figure 8: Step 2 of building G'

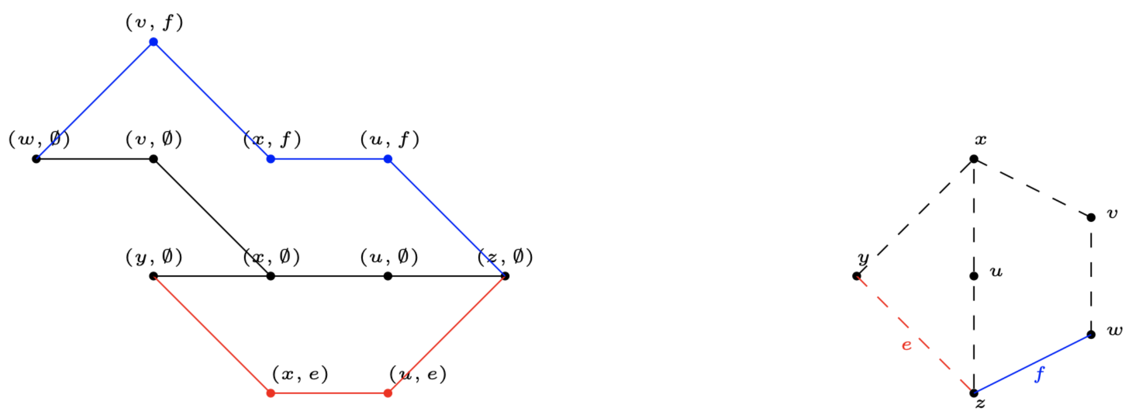

Figure 9: Step 3 of building G'

This final graph in Figure 9 is G'.

Let $\pi : X' \to X$ be the projection on the first coordinate. Then, $\pi^{-1}(x)$ is the set of all the copies of x (including the original one) that are used in the unique path in T between z_1, z_2 for some edge $e = (z_1, z_2)$. To use the Euler characteristic on G', we need a measure on X'. We define μ' on X' by

$$\mu'(A) = \int_X |\pi^{-1}(x) \cap A| \; d\mu(x)$$

Claim 1. *μ' is still finite.*

Proof of claim. Let $x \in X$. Then, $y \in \pi^{-1}(x) \iff y = (x, e)$ for some $e = (z_1, z_2)$ in G. So let's count how many paths corresponding to an edge (z_1, z_2) in G can pass through x. By our reductions, since $d_G(z_1, z_2) = 1$, we have $d_T(z_1, z_2) \leq M$ and T has maximum degree d. So we can go at most M steps "left" (from x towards z_1) and at each step, there are at most d choices, so there are at most d^M choices for z_1. Similarly, there are at most d^M choices for z_2. Thus, there are at most d^{2M} choices for the edge e and so $|\pi^{-1}(x)| \leq d^{2M}$. Then, since X is a probability space,

$$\mu'(X') = \int_X |\pi^{-1}(x) \cap X'| \, d\mu = \int_X |\pi^{-1}(x)| \, d\mu \leq \int_X d^{2M} \, d\mu = d^{2M}\mu(X) = d^{2M}$$

□

Note also that $\pi^{-1}(x)$ contains the original x so for $A \subseteq X$, $\pi^{-1}(x) \cap A = \{x\}$. So $\mu'_{|X} = \mu$.

Claim 2. *$Euler_\mu(G) = Euler_{\mu'}(G')$.*

Proof of claim. We show that $\mathcal{C}(G) + \mu'(X' \setminus X) = \mathcal{C}(G')$. If X is countable, we have a linear ordering of X. If X is uncountable, then by the Borel isomorphism theorem, X is Borel isomorphic to $2^{\mathbb{N}}$. Using the lexicographic order on $2^{\mathbb{N}}$, we get a Borel linear ordering on X. Direct the edges of G from small to large. Since G' represents the edges of G, we can direct the edges of G' as well.

Denote the out-degree of a vertex y in a graph H by $|\overrightarrow{H}_y|$. Each fake vertex has out-degree exactly one since each fake vertex is in a unique path associated with one edge. Moreover, it is clear that for $x \in X$, we have $|\overrightarrow{G}_x| = |\overrightarrow{G'}_x|$. Thus,

$$\begin{aligned}
\mathcal{C}(G') &= \int_{X'} |\overrightarrow{G'}_x| \, d\mu'(x) && \text{valency def} \\
&= \int_X |\overrightarrow{G'}_x| \, d\mu'(x) + \int_{X' \setminus X} |\overrightarrow{G'}_x| \, d\mu'(x) \\
&= \int_X |\overrightarrow{G}_x| \, d\mu(x) + \int_{X' \setminus X} 1 \, d\mu'(x) && (*) \\
&= \mathcal{C}(G) + \mu'(X' \setminus X)
\end{aligned}$$

where equation $(*)$ is obtained by the observations above and the observation after claim 1 ($\mu'_{|X} = \mu$). □

We are now ready to apply folding. We start with $G_0 = G'$ and $X_0 = X'$. Given G_n, define a relation F'_n on X_n as follows: xF'_ny iff $\pi(x) = \pi(y)$ and there is $w \in X_n$ such that xG_nw and yG_nw. This means that two vertices of G_n are related if they have the same first component (they are copies of the same original vertex) and if they have a common neighbour in G_n. Let F_n be the equivalence relation generated by F'_n.

Note that all the elements in an equivalence class of F_n have the same first component, so for all $(x, e) \in X_n$, we have $[x]_{F_n} \subseteq \pi^{-1}(\pi(x, e))$. By the proof in claim 1, we have that each equivalence class of F_n is finite. Let s be a Borel selector such that for each $x \in X$, we have $s([x]_{F_n}) = (x, \emptyset)$. Let $X_{n+1} = s(X_n)$, so $X \times \{\emptyset\} \subseteq X_{n+1}$. The graph G_{n+1} has vertex set X_{n+1} and an edge between u and v if there is $u' \in [u]_{F_n}$ and $v' \in [v]_{F_n}$ with an edge between u' and v' in G_n.

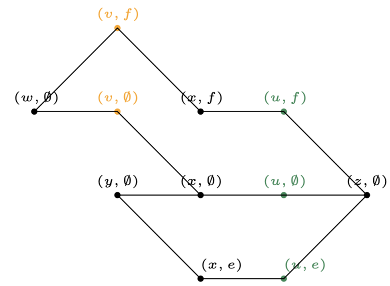

Figure 10: Step 1 of folding

Example 4. *Note that we start folding from the original vertices since that's where we will have non-trivial equivalence classes. Our non-trivial equivalence classes in $G_0 = G'$ are $\{(v, \emptyset), (v, f)\}$ in yellow and $\{(u, \emptyset), (u, e), (u, f)\}$ in green (Figure 10).*

After this first folding step, we obtain G_1 as illustrated in Figure 11. Our non-trivial equivalence class is $\{(x, \emptyset), (x, e), (x, f)\}$ in purple.

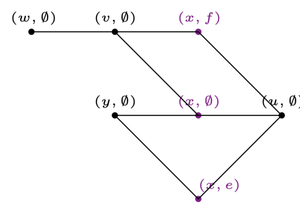

Figure 11: Step 2 of folding

After this final step of folding, we obtain G_2 as illustrated in Figure 12. Observe that $G_2 \cong T$.

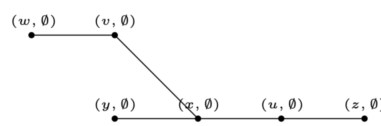

Figure 12: Step 3 of folding

Claim 3. *$Euler_(G_n) \geq Euler(G_{n+1})$.*

Proof of claim. We fix a countable Borel coloring of the equivalence classes in F_n. Let $Y_0 = X_n$ and $H_0 = G_n$. Given Y_m and H_m, we obtain H_{m+1} by folding all the classes of color m following the rules above. The vertex set is Y_{m+1}. H_m and Y_m are the intermediate graphs and spaces we obtain when folding one equivalence class of F_n at a time. We show that $Euler(H_m) \geq Euler(H_{m+1})$.

Direct the edges of H_m so that the edges with a fake vertex that gets folded as endpoint are directed away from that fake vertex. Then, by the valency definition, we have

$$\mathcal{C}(H_m) = \int_{Y_m} |\overrightarrow{H_{mx}}| \; d\mu'(x) = \int_{Y_{m+1}} |\overrightarrow{H_{mx}}| \; d\mu'(x) + \int_{Y_{m+1} \setminus Y_m} |\overrightarrow{H_{mx}}| \; d\mu'(x)$$

By how we directed the edges of H_m, we have that each fake vertex that gets folded has out-degree at least 1. Moreover, H_m stays the same on Y_{m+1}. Thus,

$$\mathcal{C}(H_m) \geq \int_{Y_{m+1}} |\overrightarrow{H_{(m+1)x}}| \, d\mu'(x) + \int_{Y_{m+1}\setminus Y_m} 1 \, d\mu'(x) = \mathcal{C}(H_{m+1}) + \mu'(Y_{m+1} \setminus Y_m)$$

Rearranging, we obtain $Euler(H_m) \geq Euler(H_{m+1})$. By induction, we get

$$Euler(G_n) = Euler(H_0) \geq \lim_{m\to\infty} Euler(H_m) = Euler(G_{n+1})$$

□

Claim 4. *The number of folding iterations is finite. In fact,* $X_{M+LM^2} = X$ *and* $G_{M+LM^2} = T$.

Proof of claim. Since the folding operation preserves the equivalence relation between the original vertices, it is enough to show that $X_{M+LM^2} = X$.

Let's look at a copy of the vertex x: $(x, e) \in X'\setminus X$, where $e = (y, y')$ is an edge in G. Then, $d_G(y, y') = 1$. By our reduction assumptions, we have $\frac{1}{M} d_T \leq d_G$, and so $d_T(y, y') \leq M$. Since (x, e) is in the path from y to y', we get $d_T(x, y) \leq M$. By our reduction assumptions, we have $d_G \leq L d_T$, and so $d_G(x, y) \leq LM$. By the triangle inequality, we obtain the following that we will analyze

$$d_{G'}((x, e), (x, \emptyset)) \leq d_{G'}((x, e), (y, \emptyset)) + d_{G'}((y, \emptyset), (x, \emptyset))$$

For the first term, recall that G' is implemented with the edges of T so

$$d_{G'}((x, e), (y, \emptyset)) = d_T(x, y) \leq M$$

For the second term, note that the path in T from y to x contains at most M edges since $d_T(x, y) \leq M$. Since $d_G \leq L d_T$, then $d_G(x, y) \leq LM$. Similarly as above, for each edge on the path from y to x in G, we have that the corresponding path in T has length at most M. Then, a path from $(y, \emptyset)$ to $(x, \emptyset)$ in G' corresponds to taking the path in T for each edge in the path from y to x in G. Thus,

$$d_{G'}((y, \emptyset), (x, \emptyset)) \leq (LM)M = LM^2$$

Combining these, we obtain

$$d_{G'}((x, e), (x, \emptyset)) \leq d_{G'}((x, e), (y, \emptyset)) + d_{G'}((y, \emptyset), (x, \emptyset)) \leq M + LM^2$$

Since T is tree and (x, e) is a copy of x, then any path from (x, e) to $(x, \emptyset)$ in G' must backtrack. This backtrack corresponds to the common vertex required for folding. Thus, (x, e) gets folded onto $(x, \emptyset)$ in at most $M + LM^2$ steps. □

Now, combining claims 2, 3 and 4, we get that

$$Euler_\mu(G) = Euler_{\mu'}(G') = Euler_{\mu'}(G_0) \geq Euler_{\mu'}(G_{M+LM^2}) = Euler_{\mu'}(T) = Euler_\mu(T)$$

where the last equality comes from the fact that $\mu'_{|X} = \mu$. Then,

$$\mathcal{C}(G) - \mu(X) \geq \mathcal{C}(T) - \mu(X) \implies \mathcal{C}(G) \geq \mathcal{C}(T)$$

as desired.

□

7 Marker Lemma

Finally, we will now see an interesting lemma that is useful in the study of CBERs and how it gives us a lower bound on the cost of some CBERs. This section follows the proofs of lemma 6.7 and corollary 21.3 of [KM04].

Definition 17 (Vanishing sequence of markers). *A vanishing sequence of markers for a CBER $\mathcal{R}$ is a decreasing sequence of Borel complete sections for $\mathcal{R}$ with empty intersection.*

Definition 18 (Aperiodic). *A CBER is aperiodic if its classes are infinite.*

Lemma 2 (Marker lemma). *Every aperiodic CBER $\mathcal{R}$ admits a vanishing sequence of markers.*

Proof. By the Borel isomorphism theorem, we can assume wlog that $\mathcal{R}$ is a CBER on $X = 2^{\mathbb{N}}$. Let $n \in \mathbb{N}$. For each $x \in X$, let $s_n(x)$ be the lexicographic least $s \in 2^n$ s.t $|\mathcal{R}(x) \cap [s]| = \infty$. In other words, it is the word s of length n with the most leading 0s s.t the equivalence class of x intersects infinitely many times the cylinder $[s]$. Such $s \in 2^n$ must exist since the cylinders $[w]$ for $w \in 2^n$ partition X, there are finitely many of them (2^n in fact) and $\mathcal{R}$ is aperiodic so each class is infinite. Note that for $x, y \in X$ s.t $\mathcal{R}(x) = \mathcal{R}(y)$, we have $s_n(x) = s_n(y)$.

Now, define $A_n \subseteq X$ as $A_n = \{x \in X : x_{|n} = s_n(x)\}$. This means A_n consists of the x that are included in their chosen cylinder $[s_n(x)]$. Since the chosen cylinders are decreasing, it is clear that the A_n are also decreasing. Each A_n is Borel since it is the intersection of a CBER and cylinders (Borel subsets).

By definition of $s_n(x)$, each A_n intersects each equivalence class infinitely many times. Let $A = \bigcap_{n \in \mathbb{N}} A_n$, which is Borel. For each class, the cylinders decrease to at most one point (and so their intersection is at most one point). Thus, A intersects each class at most once. Finally, we define our vanishing sequence of markers: $B_n = A_n \setminus A$. The intersection is empty since we removed A. It is Borel since A_n, A are Borel. They are complete sections since A_n intersect each class infinitely many times and A only intersects each class at most once. So B_n intersects each class. □

Corollary 1. *If $\mathcal{R}$ is an aperiodic CBER on X and μ is $\mathcal{R}$-invariant, then $\mathcal{C}_\mu(\mathcal{R}) \geq \mu(X)$.*

Proof. By the Marker lemma 2, let $\{B_n\}$ be a vanishing sequence of markers for $\mathcal{R}$, so each B_n is a complete section for $\mathcal{R}$. Let Φ be a graphing of $\mathcal{R}$. Then, by part 4 of lemma 1 applied to $Y = B_n$ for each $n \in \mathbb{N}$, we have $\mathcal{C}_\mu(\Phi_v) = \mu(X) - \mu(B_n)$, which gives

$$\mathcal{C}_\mu(\Phi) \geq \mathcal{C}_\mu(\Phi_v) = \mu(X) - \mu(B_n) = \mu(X \setminus B_n)$$

Since the B_n are vanishing, we have

$$X = X \setminus \bigcap_{n \in \mathbb{N}} B_n = \bigcup_{n \in \mathbb{N}} (X \setminus B_n)$$

Since the B_n are decreasing, the $(X \setminus B_n)$ are increasing. By monotone convergence, we obtain

$$\lim_{n \to \infty} \mu(X \setminus B_n) = \mu(X)$$

Hence, $\mathcal{C}_\mu(\Phi) \geq \mu(X)$. □

8 Conclusion

In this paper, we have explored the concept of cost for graphs, equivalence relations and groups. This invariant quantifies their complexity and how efficiently we can express them. Because the definition of cost for CBERs and groups is often difficult to compute, we investigated alternative methods, namely the induction formula and the Fundamental Theorem of Cost to provide other means of computation. Finally, we proved the Marker's lemma, which is useful in the study of CBERs, and used it to provide a lower bound for aperiodic CBERs.

9 Acknowledgements

I would like to thank Owen Rodgers, my mentor,. for his guidance, support and passion in this Directed Reading Program. I would also like to thank Sasha Bell, Catherine Fontaine and Hazem Hassan for organizing the Winter 2026 DRP.

10 References

[Gab] Damien Gaboriau. *Around the orbit equivalence theory, measure equivalence, cost and ℓ^2 Betti numbers.* URL: https://perso.ens-lyon.fr/gaboriau/Travaux-Publi/ME-Cost-L2-Lectures/ME-Cost-L2-lectures.pdf.

[Gab00] Damien Gaboriau. "Coût des relations d'équivalence et des groupes". fr. In: *Invent. Math.* 139.1 (Jan. 2000), pp. 41–98.

[KM04] Alexander S. Kechris and Benjamin D. Miller. *Topics in orbit equivalence.* Springer, 2004.

[Zom20] Jenna Zomback. *"Mercuriale de Groupes et de Relations" by Damien Gaboriau.* 2020. URL: https://drive.google.com/file/d/14kZM8O7OE7tRiCEcKQ5yE5eMYdUsG7c-/view.