Point processes on locally compact groups and their cost

by
Samuel Anthony Mellick

Submitted to
Central European University
Department of Mathematics and its Applications

In partial fulfillment of the requirements
for the degree of Doctor of Philosophy in Mathematics

Supervisor: Miklós Abért
Alfréd Rényi Institute of Mathematics

CENTRAL EUROPEAN UNIVERSITY

Budapest, Hungary
September 2019

Doctoral thesis
Samuel Anthony Mellick
Central European University, Budapest, Hungary
Printed in Hungary

Declaration

I, the undersigned Samuel Anthony Mellick, candidate for the degree of Doctor of Philosophy at the Central European University Department of Mathematics and its Applications, declare herewith that the present thesis is based on my research and only such external information as properly credited in notes and bibliography.

I declare that no unidentified and illegitimate use was made of the work of others, and no part of this thesis infringes on any person's or institution's copyright. I also declare that no part of this thesis has been submitted to any other institution of higher education for an academic degree.

Budapest, Hungary, September 2019

Samuel Anthony Mellick

Acknowledgements

Thanks to my supervisor Miklós Abért for exposing me to a new area of maths and giving me the opportunity to explore it.

Thanks to my fellow student Lászlo Márton Tóth for companionship and all his help navigating the bureaucracy.

Thanks to Mikołaj Fraczyk for putting up with my endless distractions of him and for listening to my half-baked, quarter-baked, and (on occasion) eighth-baked ideas. I tried to get the ogonek but it didn't work, sorry.

Thanks to Alessandro Carderi for enduring a probability theory lecture from me and explaining the cross-section side of things.

Thanks to my family and my friends Andy, Michelle, Brooke, Matt, and Anna for their support.

Contents

Introduction

The topic of this thesis is the study of *point processes* from the measured group theory perspective. The groups will be locally compact and second countable (lcsc) and *nondiscrete*. A point process on such a group G is then a *random discrete subset* $\Pi \subset G$. We will use the general theory of such processes, but restrict our attention to such point processes which are *distributionally invariant* in the sense that $g\Pi \overset{d}{=} \Pi$ for all $g \in G$.

One of the main characters in the thesis will be the *Poisson point process* on G. This is a model that is, in some sense, "maximally random". It is to nondiscrete groups G as *Bernoulli percolation* is to discrete groups Γ. We will also study a model called *the IID Poisson point process*, which should be viewed as the analogue of the full Bernoulli shift $\Gamma \curvearrowright [0,1]^\Gamma$.

Our entry point in the study of invariant point processes is the following fact that we prove:

Fact. Associated to a (finite intensity) invariant point process Π with distribution μ is a *probability measure preserving* (pmp) measure μ_0 on a *countable Borel equivalence relation* (cber) $(\mathbb{M}_0, \mathcal{R})$.

We refer to this object as the *Palm equivalence relation* of Π.

The utility of this statement is that certain questions about factor constructions that "live on" the point process correspond exactly to *known* concepts on the pmp cber. For example, let $\mathscr{G}$ denote a *factor graph* of Π: this is a graph whose vertex set is Π and is constructed *deterministically* in an *equivariant* way, so that $\mathscr{G}(g\Pi) = g\mathscr{G}(\Pi)$. Under the correspondence principle mentioned above, *connected* factor graphs correspond exactly to *graphings* of the associated Palm equivalence relation.

Questions about deterministic and factor of IID graphs on a point process are incredibly natural from the probability theory perspective. Building on work of Holroyd and Peres [HP03] and later Timár [Tim04], we prove:

Theorem 1. Let Π be an essentially free and ergodic point process on G. Then Π admits a factor graph which is isomorphic to $\mathbb{Z}$ if and only if G is amenable. In that case, Π admits factor graphs of the form $\mathrm{Cay}(A, S)$, where $A = \langle S \rangle$ is any finitely generated amenable group.

This extends their work to the maximum possible level of generality, but this was not our goal per se. The point is that once one is aware of the cber structure associated to a point process, the above question becomes trivial from existing results.

We observe that the IID Poisson point process is *cocycle superrigid for discrete targets*, and use some point process theory to establish the following:

Theorem 2. Let G be a group with Kazhdan's property (T) and no compact normal subgroups. Then the Poisson point process Π on G admits *no* connected factor graph of the form $\mathrm{Cay}(\Gamma, S)$ for *any* discrete group $\Gamma = \langle S \rangle$.

We choose to state the theorem in the above form, but it has an equivalent statement in terms of the associated Palm equivalence relation of Π: it cannot be *freely* generated by any action of a discrete group Γ. The first examples of pmp cbers with this property were produced by Furman in [Fur99], see also Thomas [Tho03].

Invariant point processes form a rich and interesting class of probability measure preserving (pmp) actions of lcsc groups. In fact, in a certain sense they are *the only* examples:

Theorem 3. Every essentially free pmp action of an lcsc groups is abstractly isomorphic to an invariant point process.

We are particularly interested in a numerical invariant of invariant point processes called *cost*. This definition was suggested by Miklós Abért, inspired by work of Gaboriau and Levitt. The cost of a point process is the "cheapest" way to wire it up as a connected graph. It turns out that this notion is exactly equivalent to the existing concept of cost for the pmp cber we associate to point processes.

Using the aforementioned isomorphism theorem, one can define cost for any free pmp action of an lcsc group. We prove that this is well-defined (that is, an isomorphism invariant) by proving a version of Gaboriau's theorem on the cost of a "complete section", but completely in the point process language.

It has been known in the community that one can define a consistent notion of cost for free pmp actions of *unimodular* lcsc groups, by first fixing[1] a Haar measure λ. However, until recently (see [Car18]) this was not done, at least publicly. The reason for this is that whilst one can *define* cost, one couldn't actually *prove* anything about it.

With the point process framework however, we are able to make a first step:

Theorem 4. Amongst all free point processes on a group G, the Poisson point process has *maximal cost*.

As the name suggests, cost is a *nonnegative* invariant. In fact, by a quirk of normalisation it is in fact always at least one. Thus the previous theorem gives a strategy to show that *all* free point processes on a group have *the same* cost: prove that the Poisson point process has cost one. Such a group is said to have *fixed price one*.

To that end, we prove the following theorem:

Theorem 5. All groups of the form $G \times \mathbb{Z}$ have fixed price one.

We also extend this to groups of the form $G \times \mathbb{R}$ and sketch how it will go for groups more generally containing a noncompact amenable normal subgroup.

Cost has connections with other interesting asymptotic invariants also. We prove the following connection with rank gradient:

Theorem 6. Let (Γ_n) be a *Farber sequence* of *cocompact* lattices in a group G. Suppose further that the sequence is *uniformly uniformly*[2] *discrete*. Then

$$\limsup_{n\to\infty} \frac{d(\Gamma_n) - 1}{\operatorname{covol} \Gamma_n} \leq \operatorname{cost}(\text{Poisson on } G) - 1,$$

where $d(\Gamma)$ denotes the *rank* of of Γ, that is, the minimal size of a generating set of Γ.

[1] This is reminiscent of Petersen's work [Pet13] on the related topic of L^2 Betti numbers.

[2] Not a typo.

The above theorem was essentially proved in an independent work by Carderi [Car18], but without the connection to the Poisson point process (he bounds the rank gradient in terms of the cost of an ultraproduct of actions).

The structure of the thesis is as follows:

In Chapter 1, We give a basic overview of point processes, with the goal of educating the uninitiated. The main tasks here are to establish notation and to construct the associated *Palm equivalence relation* of a point process and investigate its properties. We give a few applications of this equivalence relation.

In Chapter 2, We collect various known facts about point processes as they pertain to the concept of *weak convergence*. This is to serve as a convenient reference for the reader.

In Chapter 3, We prove an analogue of the Abért-Weiss theorem for point processes, and then the aforementioned theorems on cost. This makes use of a technique we call *weak factoring*, which is inspired by a notion called *weak containment*. The document concludes with a list of questions about how weak factoring might connect to weak containment, and explains the connections between point process theory as we've discussed it and the notion of *cross-section equivalence relations*.

The background material I used to teach myself about point processes is Kingman's book [Kin93] and the two volumes by Daley and Vere-Jones [VJ03] [DVJ07], the latter mainly as a reference. In the course of writing the thesis two excellent resources became available: the book of Last and Penrose [LP18], and some lecture notes by Blaszczyszyn [Bla17].

Chapter 1

Point processes and the Palm equivalence relation

1.1 Point process basics and setting the scene

Let (Z, d) denote a complete and separable metric space (a csms). A *point process on Z* is a random discrete subset of Z. We will also study random discrete subsets of Z that are *marked* by elements of an additional csms Ξ. Typically Ξ will be a finite set that we think of as colours.

Definition 1. The *configuration space* of Z is

$$\mathbb{M}(Z) = \{\omega \subset Z \mid \omega \text{ is discrete}\},$$

and the Ξ-*marked configuration space* of Z is

$$\Xi^{\mathbb{M}}(Z) = \{\omega \subset Z \times \Xi \mid \omega \text{ is discrete, and if } (g, \xi) \in \omega \text{ and } (g, \xi') \in \omega \text{ then } \xi = \xi'\}.$$

Note that $\Xi^{\mathbb{M}}(Z) \subset \mathbb{M}(Z \times \Xi)$. We think of a Ξ-marked configuration $\omega \in \Xi^{\mathbb{M}}(Z)$ as a discrete subset of Z with labels on each of the points (whereas a typical element of $\mathbb{M}(Z \times \Xi)$ is a discrete subset where each point has possibly multiple marks).

If $\omega \in \Xi^{\mathbb{M}}(Z)$ is a marked configuration, then we will write ω_z for the unique element of Ξ such that $(z, \omega_z) \in \omega$.

The Borel structure on configuration spaces is exactly such that the following *point counting functions* are measurable. Let $U \subseteq Z$ be a Borel set. It induces a function $N_U : \mathbb{M}(Z) \to \mathbb{N}_0 \cup \{\infty\}$ given by

$$N_U(\omega) = |\omega \cap U|.$$

We will primarily be interested in point processes defined on locally compact and second countable (lcsc) groups G. Such groups admit a unique (up to scaling) Haar measure λ, we fix such a choice. Recall:

Theorem 7 (Struble's theorem, see Theorem 2.B.4 of [CdlH16])**.** *Let G be a locally compact topological group. Then G is second countable* if and only if *it admits a proper*[1] *left-invariant metric.*

[1]Recall that a metric is *proper* if closed balls are compact.

Such a metric is unique up to coarse equivalence (bilipschitz if the group is compactly generated). We fix d to be any such metric.

We mostly consider the configuration space of a fixed group G. So out of notational convenience let us write $\mathbb{M} = \mathbb{M}(G)$ and $\Xi^{\mathbb{M}} = \Xi^{\mathbb{M}}(G)$. The latter here is an abuse of notation: formally $\Xi^{\mathbb{M}}$ ought to denote the set of functions from $\mathbb{M}$ to Ξ, but instead we are using it to denote the set of functions from *elements* of $\mathbb{M}$ to Ξ.

Note that the marked and unmarked configuration spaces of G are Borel G-spaces. To spell this out, $G \curvearrowright \mathbb{M}$ by $g \cdot \omega = g\omega$ and $G \curvearrowright \Xi^{\mathbb{M}}$ by

$$g \cdot \omega = \{(gx, \xi) \in G \times \Xi \mid (g, \xi) \in \omega\}.$$

Definition 2. A *point process* on G is a $\mathbb{M}(G)$-valued random variable $\Pi : (\Omega, \mathbb{P}) \to \mathbb{M}(G)$. Its *law* or *distribution* μ_Π is the pushforward measure $\Pi_*(\mathbb{P})$ on $\mathbb{M}(G)$. It is *invariant* if its law is an invariant probability measure for the action $G \curvearrowright \mathbb{M}(G)$.

The associated *point process action* of an invariant point process Π is $G \curvearrowright (\mathbb{M}(G), \mu_\Pi)$.

Some remarks and caveats are in order:

- Point processes which are not invariant are very much of interest, but the only examples which we will consider will be so-called "Palm point processes", to be defined later. Thus unless explicitly prefaced by the word "Palm", one ought to interpret "point process" as "invariant point process".
- Speaking properly, we are discussing *simple* point processes, that is, those where each point has multiplicity one. We will discuss this more later.
- Ξ-marked point processes are defined similarly, with $\Xi^{\mathbb{M}}$ taking the place of $\mathbb{M}$. There isn't much difference between marked point processes and unmarked ones for our purposes (it's just a case of which is more convenient for the particular problem at hand). Thus "point process" might also mean "marked point process".
- One could certainly define point processes on a discrete group, but this is better known as percolation theory. We are trying to move beyond that, so we will almost always implicitly assume G is nondiscrete.
- The other case of interest is $\mathrm{Isom}(S)$-invariant point processes on S, where S is a Riemannian symmetric space. For instance, one would consider isometry invariant point processes on Euclidean space $\mathbb{R}^n$ or hyperbolic space $\mathbb{H}^n$.
- Our interest in point processes is almost exclusively *as actions*. We will therefore rarely distinguish between a point process proper and its distribution. Thus we may use expressions like "suppose μ is a point process" to mean "suppose μ is the distribution of some point process".
- The configuration space of any Polish space will be Polish, so the probability theory of point processes on such spaces is well behaved. We will have more to say on the metric structure of the configuration space later.

Definition 3. The *intensity* of a point process μ is

$$\mathrm{intensity}(\mu) = \frac{1}{\lambda(U)} \mathbb{E}_\mu \left[N_U \right],$$

where $U \subset G$ is any Borel set of positive (but finite) Haar measure, and $N_U(\omega) = |\omega \cap U|$ is its point counting function.

To see that the intensity is well-defined (that is, does not depend on our choice of U), observe that the function $U \mapsto \mathbb{E}_\mu[N_U]$ defines a Borel measure on G which inherits invariance from the shift invarance of μ. So by uniqueness of Haar measure, it is some scaling of our fixed Haar meausure λ – the intensity is exactly this multiplier. We also see that whilst the intensity depends on our choie of Haar measure, it scales linearly with it.

Note that a point process has intensity zero if and only if it is empty almost surely.

1.1.1 Examples of point processes

Example 1 (Lattice shifts)**.** Let $\Gamma < G$ be a *lattice*, that is, a discrete subgroup that admits an invariant probability measure ν for the action $G \curvearrowright G/\Gamma$. The natural map $\mathbb{M}(G/\Gamma) \to \mathbb{M}(G)$ given by

$$\omega \mapsto \bigcup_{a\Gamma \in \omega} a\Gamma$$

is left-equivariant, and hence maps invariant point processes on G/Γ to invariant point processes on G. In particular, we have the *lattice shift*, given by choosing a ν-random point $a\Gamma$.

Example 2 (**Induction from a lattice**)**.** Now suppose one also has a pmp action $\Gamma \curvearrowright (X, \mu)$. It is possible to *induce* this to a pmp action of G on $G/\Gamma \times X$. This can be described as an X-marked point process on G. To do this, fix a fundamental domain $\mathscr{F} \subset G$ for Γ. Choose $f \in \mathscr{F}$ uniformly at random, and independently choose a μ-random point $x \in X$. Let

$$\Pi = \{(f\gamma, \gamma \cdot x) \in G \times X \mid \gamma \in \Gamma\}.$$

Then Π is a G-invariant X-marked point process.

So point processes can be thought of as "generalised lattices". For this theory to be nontrivial then, one must have examples which are genuinely non-lattice. We will define the most fundamental example of a point process after the following `PROB101` refresher:

Recall that a random integer N is *Poisson distributed with parameter* $t > 0$ if $\mathbb{P}[N = k] = \exp(-t)\frac{t^k}{k!}$. We write $N \sim \texttt{Pois}(t)$ to denote this. It is convenient to extend this definition to $t = 0$ and $t = \infty$ by declaring $N \sim \texttt{Pois}(0)$ when $N = 0$ almost surely and $N \sim \texttt{Pois}(\infty)$ when $N = \infty$ almost surely.

Definition 4. Let X be a complete and separable metric space equipped with a non-atomic Borel measure λ.

A point process Π on X is *Poisson with intensity* $t > 0$ if it satisfies the following two properties:

(Poisson point counts) for all $U \subseteq G$ Borel, $N_U(\Pi)$ is a Poisson distributed random variable with parameter $t\lambda(U)$, and

(Total independence) for all $U, V \subseteq G$ *disjoint* Borel sets, the random variables $N_U(\Pi)$ and $N_V(\Pi)$ are *independent.*

For reasons that should not be immediately apparent, both of the above defining properties are equivalent. We will write $\mathscr{P}_t$ for the distribution of such a random variable, or simply $\mathscr{P}$ if the intensity is understood.

We think of the Poisson point process as a completely random scattering of points in the group. It is an analogue of Bernoulli site percolation for a continuous space.

We construct the process (somewhat) explicitly. Partition G into disjoint Borel sets $U_1, U_2, \ldots$ of positive but finite volume. For each of these, independently sample from a Poisson distribution with parameter $t\lambda(U_i)$. Place that number of points in the corresponding U_i (uniformly at random).

For proofs of basic properties of the Poisson point process (such as the fact that it does not depend on the partition chosen above), see the first five (breezy) chapters of Kingman's book [Kin93].

We now describe what certain sampling rules mean in terms of the implied measure space constructions, so that one can rest assured that the above description of the Poisson point process can really be done rigorously.

Let $U \subseteq G$ be a Borel subset with $0 < \lambda(U) < \infty$. When we say *a random point in U*, all we mean is some U-valued random variable whose distribution is the *probability* measure $\frac{\lambda(\bullet \cap U)}{\lambda(U)}$.

When we say *take n random (unordered) points in U*, what we mean is that you should take the map $F : U^n \to \mathbb{M}(U)$ given by

$$F(u_1, u_2, \ldots, u_n) \mapsto \{u_1, u_2, \ldots, u_n\}$$

and *pushforward* under it the measure $\left(\frac{\lambda(\bullet \cap U)}{\lambda(U)}\right)^{\otimes n}$. When we say *take N random points in U, where $N \sim$ **Pois**$(t\lambda(U))$*, we mean you should take the map $F : \mathbb{N}_0 \times U^{\mathbb{N}}$ given by

$$F(n; u_1, u_2, \ldots) \mapsto \{u_1, u_2, \ldots, u_n\}$$

and pushforward under it the measure $\mathcal{L}(N) \otimes \left(\frac{\lambda(\bullet \cap U)}{\lambda(U)}\right)^{\otimes \mathbb{N}}$, where

$$\mathcal{L}(N) = (\mathbb{P}[N = k])_{k \in \mathbb{N}_0} = (\exp(-t\lambda(U))\frac{(t\lambda(U))^k}{k!}$$

denotes the distribution of N.

At last, we explicitly describe what the Poisson point process will look like: decompose G as a disjoint union of positive (but finite measure) sets $G = \bigsqcup_i U_i$. On each of these, independently sample N_i points uniformly at random, where $N_i \sim \texttt{Pois}(t\lambda(U_i))$. Let $F^i : \mathbb{N}_0 \times U^{\mathbb{N}}$ be the maps that implement that. Then consider the map

$$\prod_i (\mathbb{N}_0 \times U^{\mathbb{N}}) \to \mathbb{M}(G), (n^i; u_1^i, u_2^i, \ldots)_i \mapsto \bigcup_i F^i(n^i; u_1^i, u_2^i, \ldots, u_{n^i}^i),$$

and pushforward under it the measure $\bigotimes_i \mathcal{L}(N_i) \otimes \left(\frac{\lambda(\bullet \cap U_i)}{\lambda(U_i)}\right)^{\otimes \mathbb{N}}$.

This rather distasteful construction now complete, the reader ought to take a break and imbibe a palette cleanser. We recommend a pint of vodka.

Definition 5. A pmp action $G \curvearrowright (X, \mu)$ is *ergodic* if for every G-invariant subset $A \subseteq X$, we have $\mu(A) = 0$ or $\mu(A) = 1$.

The action is *mixing* if for all $A, B \subseteq (X, \mu)$ we have

$$\lim_{g \to \infty} \mu(gA \cap B) = \mu(A)\mu(B).$$

Proposition 1. The Poisson point process actions $G \curvearrowright (\mathbb{M}, \mathscr{P}_t)$ on a noncompact group G are essentially free and ergodic (in fact, mixing).

A proof of freeness that is readily adaptable to our setting can be found as Proposition 2.7 of [ABB⁺17]. For ergodicity and mixing, see the proof of the discrete case in Proposition 7.3 of the Lyons-Peres book [LP16]. It directly adapts, once one knows the required cylinder sets exist.

Although the subscript t suggests that the Poisson point processes form continuum family of actions, this is not always the case:

Theorem 8 (Ornstein-Weiss)**.** Let G be an amenable group which is not a countable union of compact subgroups. Then the Poisson point process actions $G \curvearrowright (\mathbb{M}, \mathscr{P}_t)$ are all isomorphic.

Question 1. Can the Poisson point processes of different intensities be nonisomorphic over a nonamenable group G?

The following definition uses notation that does not appear in the literature (the object of course does, but there does not appear to be a symbolic representation for it):

Definition 6. If Π is a point process, then its *IID version* is the $[0,1]$-marked point process $[0,1]^\Pi$ with the property that conditional on its set of points, its labels are independent and IID Unif$[0,1]$ distributed. If μ is the law of Π, then we will write $[0,1]^\mu$ for the law of $[0,1]^\Pi$.

One can define the IID of a point process over spaces other than $[0,1]$ (for instance, $[n] = \{1, 2, \ldots, n\}$ with the counting measure), but we will only use the full IID.

Remark 1. This process really exists. Fix an *enumeration rule* $\texttt{enum} : \mathbb{M} \to \mathbb{N}^\mathbb{M}$. This is a measurable (but *necessarily not* G-equivariant) map with the property that $\texttt{enum}(\omega)$ is an enumeration $(\texttt{enum}(\omega)_1, \texttt{enum}(\omega)_2 \ldots)$ of ω.

Now consider the probability space

$$(\mathbb{M}, \mu) \otimes \prod_{n \in \mathbb{N}} ([0,1], \text{Leb}).$$

and the following labelling map L from it into $[0,1]^\mathbb{M}$

$$L(\omega, (\xi_1, \xi_2, \ldots)) = \{(\texttt{enum}(\omega)_n, \xi_n) \in G \times [0,1] \mid n \in \mathbb{N}\}.$$

The IID version of μ is the pushforward of $\mu \otimes \text{Leb}^{\otimes \mathbb{N}}$ under this map. Note that this map is *not equivariant*, but nevertheless the pushfoward is *distributionally* invariant, and thus defines an invariant point process. For further details see section 5.2 of [LP18].

Remark 2. As we've mentioned, $[0,1]$-marked point processes on G are particular examples of point processes on $G \times [0,1]$. One can show (see Theorem 5.6 of [LP18]) that the Poisson point process on $G \times [0,1]$ with respect to the product measure $\lambda \otimes \texttt{Leb}$ is just the IID version of the Poisson point process on G, a fact which we will make use of later.

1.1.2 Characterising noncompactness through point processes

Proposition 2. Let Π be a point process on a noncompact group. Then $|\Pi| = 0$ or ∞ almost surely.

Point processes on compact groups are always finite (as is any discrete subset of a compact space).

We will give two proofs: the first I became aware of via a MathOverflow thread [nh], the second is the usual proof for $\mathbb{R}^n$.

Proof one. Suppose Π is finite but non-empty with positive probability. Then conditional on Π being finite but non-empty, a point of Π chosen uniformly at random will be equidistributed on G – that is, G admits a *finite* Haar measure. It is known that a group has a finite Haar measure if and only if it is compact, finishing the proof.

Let us be more precise. Denote by μ the law of Π, and $[0,1]^\mu$ the law of $[0,1]^\Pi$.

Consider the map $F : [0,1]^{\mathbb{M}} \to G$ given by

$$F(\omega) = g, \text{ where the label of } g \text{ is } \textit{least} \text{ amongst all elements of } \omega.$$

This map is defined on the set $E = \{\omega \in [0,1]^{\mathbb{M}} \mid 0 < |\omega| < \infty\}$, and is equivariant.

If Π is finite but non-empty with positive probability, then the same is true of $[0,1]^\Pi$ – that is, E has positive measure with respect to $[0,1]^\mu$. The pushforward of $[0,1]^\mu$ *restricted* to E under F therefore defines a *finite* Borel measure on G. This implies G is compact, a contradiction. □

Proof two. Suppose $k \in \mathbb{N}$ is such that $\mathbb{P}[|\Pi| = k] > 0$. Then let Π_k denote Π conditioned to contain exactly k points. This defines a shift invariant process.

Let $U \subseteq G$ be a bounded set of unit Haar measure. By noncompactness of G we can find an infinite sequence $g_1, g_2, \ldots$ such that all translates $\{g_n U\}$ are disjoint. Then by shift invariance $\mathbb{E}[|\Pi_k \cap g_n U|] = \mathbb{E}[|\Pi_k \cap U|]$ for all n. So

$$k = \mathbb{E}[|\Pi_k|] \geq \mathbb{E}\left[\left|\bigcup_n \Pi_k \cap g_n U\right|\right] \geq \sum_n \mathbb{E}\left[|\Pi_k \cap g_n U|\right] = \infty.$$

This is a contradiction, so $\mathbb{P}[|\Pi| = k] = 0$ for every k, as desired. □

The second proof can be extended to show that *invariant random measures* on a noncompact space have total mass 0 or ∞ almost surely, see Proposition 12.1.VI of [DVJ07].

Remark 3. The first proof is slick and appealing for the probabilistically minded, but both proofs essentially use the same contradiction: that in a noncompact group, one can find infinitely many disjoint translates of a fixed bounded set.

Note that the Poisson point process is "fully random": a typical sample will contain any local behaviour that you want (or want to avoid). Moreover, it will occur infinitely often. More formally:

Proposition 3. Let Π be a sample from the Poisson point process. Then its orbit $G \cdot \Pi$ in $\mathbb{M}(G)$ is dense.

This statement is not used elewhere in the work, and will not make sense until Section 2, so feel free to skip the proof.

Proof. The idea of the proof is as follows: any finite configuration will (approximately) appear in a fixed region R in any sample of a Poisson point process with positive probability. If we look for this finite configuration in an infinite collection of *disjoint* translates of R, then by independence it will almost surely occur (in fact, it will occur infinitely often). By taking enough finite configurations and enough infinite collections of disjoint regions, one can ensure that any finite configuration will approximately appear *somewhere* in the Poisson. That is exactly saying its orbit is dense.

Choose a dense sequence (ω_n) of configurations in $\mathbb{M}(G)$. For each ω_n, fix a sequence (r_{n_k}) of radii such that if $\pi_k \in \mathbb{M}(G)$ is a sequence of configurations, then

$$\pi_k \text{ converges to } \omega_n \text{ if and only if for every } k \in \mathbb{N}, d_{\text{prok}}(\pi_k|_{B(0,r_{n_k})}, \omega_n|_{B(0,r_{n_k})}) \to 0,$$

where $B(0, r_{n_k})$ denotes the radius r_{n_k} ball about the identity 0 in G.

For each n and k, choose an infinite sequence of points $x_{n,k,i}$ such that all of the balls $B(x_{n,k,i}, r_{n_k})$ are disjoint. This is possible by noncompactness.

At last, consider the *events*

$$E_{n,k,i} = \left\{ d_{\text{prok}} \left(x_{n,k,i}^{-1} \Pi|_{B(x_{n,k,i},r_{n_k})}, \omega_n|_{B(0,r_{n_k})} \right) < \frac{1}{n_k} \right\}.$$

Then $\mathbb{P}[\forall n, k\, \Pi \in E_{n,k,i} \text{ infinitely often}] = 1$ by Borel-Cantelli, and on this event Π has a dense orbit in $\mathbb{M}(G)$. □

In particular, a Poisson point process is unlatticelike in the sense that it will contain arbitrarily large "holes" and very small regions completely crammed with points. The following model rectifies this:

Example 3 (The Poisson net)**.** Let $\varepsilon > 0$. From an independent sequence $\Pi_1, \Pi_2, \ldots$ of unit intensity[2], we construct a point process Π in the following way:

$$\Pi = \bigcup_n \{g \in \Pi_n \mid \text{ for all } i \leq n, d(g, \Pi_i \setminus \{g\}) > \varepsilon\}.$$

That is, we retain all points of Π_1 that are ε-separated, and then add points from the remaining Poissons that retain this ε-separation property.

The resulting process is thus *uniformly discrete*[3]. Moreover, it is *coarsely dense* in the sense that every point of G is distance at most 2ε from Π.

Thus Π is a randomly constructed ε-net (in the sense of metric geometry), and we will refer to it as *the Poisson net.*

Question 2. Is the Poisson net free? It seems that the answer must be yes, but I have not proved this. Note that the proof must use nondiscreteness in an essential way (as the analogous model with $\varepsilon < 1$ on any discrete group Γ results in Γ itself).

The Poisson net is certainly ergodic, as it is constructed as a factor of the ergodic system $G \curvearrowright (\mathbb{M}, \mathscr{P})^{\otimes \mathbb{N}}$, where G acts diagonally. One can prove that this latter system is ergodic as a consequence of $G \curvearrowright (\mathbb{M}, \mathscr{P})$ being mixing, or one can construct the Poisson net as a factor of the clearly ergodic system $G \curvearrowright ([0,1]^{\mathbb{M}}, [0,1]^{\mathscr{P}})$.

Note that the IID of the Poisson net (indeed, the IID of any point process) is automatically free, as the labels break any possible symmetries. Thus if freeness is essential, one can simply use the IID of the Poisson net. This is enough for some purposes.

[2]It's not clear to what extent the model depends on the choices of intensities. For our purposes, all we care about is that such a model exists at all.

[3]Probabilists would use the term "hard-core"

Remark 4. The Poisson net can serve as a measurable substitute for a cocompact lattice in groups that do not have such lattices.

Example 4 (Periodic Poissons)**.** Suppose $\Gamma < G$ is a *discrete* subgroup. Then we may identify $\mathbb{M}(G/\Gamma)$ with the subspace of $\mathbb{M}(G)$ which is *right invariant* under multiplication by Γ, that is, configurations $\omega \in \mathbb{M}(G)$ such that $\omega\Gamma = \omega$. We will refer to such configurations as Γ-*periodic.*

Note that G/Γ carries a G-invariant measure $\lambda_{G/\Gamma}$ (this is true for any unimodular subgroup of G, see Chapter VI of [KEK05]). We will be interested in some subgroups that are not necessarily of finite covolume, so we choose to normalise $\lambda_{G/\Gamma}$ in the following way: let $\mathscr{F} \subseteq G$ be a Borel *fundamental domain* for Γ, so that

$$G = \bigsqcup_{\gamma \in \Gamma} \mathscr{F}\gamma.$$

Write $\pi : G \to G/\Gamma$ for the projection $a \mapsto a\Gamma$. Then for $B \subseteq G/\Gamma$, we set

$$\lambda_{G/\Gamma}(B) = \lambda(\pi^{-1}(B) \cap \mathscr{F}).$$

In particular, the total mass $\lambda_{G/\Gamma}(G/\Gamma)$ is $\operatorname{covol}\Gamma$.

Definition 7. The Γ-*periodic Poisson* is the Poisson point process on G/Γ sampled according to $\lambda_{G/\Gamma}$ and viewed as an invariant point process Π_Γ on G. We denote its law by $\mathscr{P}_{G/\Gamma}$.

We can explicitly describe the process as follows: let $\Pi_{\mathscr{F}}$ denote a Poisson point process on $\mathscr{F}$ sampled according to $\lambda\big|_{\mathscr{F}}$. Then $\Pi_\Gamma = \Pi_{\mathscr{F}}\Gamma$ is the result of extending $\Pi_{\mathscr{F}}$ to be Γ-periodic, hence the name.

By our choice of normalisation, intensity $\Pi_\Gamma = 1$ for any $\Gamma < G$.

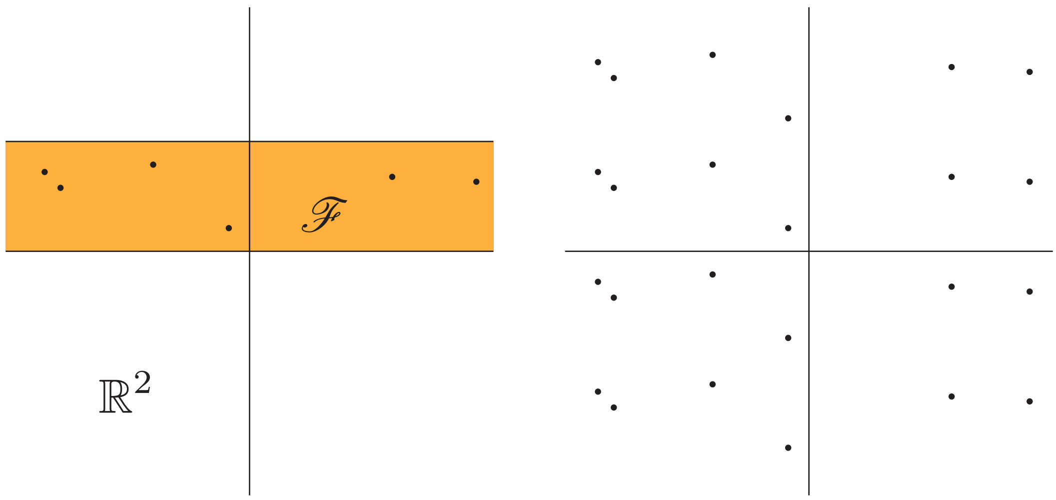

Figure 1.1: An example periodic Poisson on $\mathbb{R}^2$. On the left we sample a Poisson in the fundamental domain, on the right we see its periodic extension.

Now suppose $\Gamma < G$ is a *lattice*, so $\operatorname{covol}\Gamma < \infty$. We can explicitly describe the point counts $N_C(\Pi_\Gamma)$ for arbitrary $C \subseteq G$ as follows: denote by $a_1\Gamma, a_2\Gamma, \ldots$ an IID sequence of cosets $a_i \in G/\Gamma$ chosen according to the G-invariant probability measure on G/Γ. Let $M \sim \texttt{Pois}(\operatorname{covol}\Gamma)$ be independent of that. Then

$$N_C(\Pi_\Gamma) = \sum_{i=1}^{M} |a_i\Gamma \cap C|.$$

Distributions of this form[4] are known as *compound Poisson distributions*.

Remark 5. If $\Gamma < G$ is a lattice, then the Γ-periodic Poisson Π_Γ is very much *not* ergodic. It ergodically decomposes as a countable sum of the process "sample n independent copies of the lattice shift $G \curvearrowright G/\Gamma$ and take their union". Note in particular that it is the *empty process* with positive probability.

Remark 6. Let $\Gamma < G$ be a lattice, and Π_Γ the associated periodic Poisson. We wish to stress the difference between two different actions: the IID of Π_Γ (that is, $[0,1]^{\Pi_\Gamma}$), and the Γ-periodic IID Poisson.

By the latter, we mean you repeat the construction from earlier, but sample from the Poisson point process on $(G/\Gamma \times [0,1], \operatorname{covol} \otimes \texttt{Leb})$. This yields a $[0,1]$-labelled point process Υ on G with the following property: if $g \in \Upsilon$, then $g\gamma \in \Upsilon$ for all $\gamma \in \Gamma$, and $\Upsilon_g = \Upsilon_{g\gamma}$, where $\Upsilon_g \in [0,1]$ denotes the label of g.

Definition 8. A *point process factor map* is a G-equivariant and measurable map $\Phi : \mathbb{M} \to \mathbb{M}$. If μ is a point process and Φ is only defined μ almost everywhere, then we will call it a *μ factor map.*

We will be interested in two monotonicity conditions:

- if $\Phi(\omega) \subseteq \omega$ for all $\omega \in \mathbb{M}$, we will call Φ a *thinning* (and usually denote it by θ), and
- if $\Phi(\omega) \supseteq \omega$ for all $\omega \in \mathbb{M}$, we will call Φ a *thickening* (and usually denote it by Θ).

We use the same terms for marked point processes as well.

Remark 7. There are *two* possible ways to interpret the above monotonicity conditions for a Ξ-marked point process, depending on what you want to do with the mark space. One can consider

$$\Phi : \Xi^{\mathbb{M}} \to \Xi^{\mathbb{M}}, \text{ or } \Phi : \Xi^{\mathbb{M}} \to \mathbb{M}.$$

In the former case, the definition above works verbatim. In the latter case, one should interpret a statement like "$\omega \subseteq \Phi(\omega)$" as "$\omega$ is contained in the underlying set $\pi(\Phi(\omega))$ of $\Phi(\omega)$, where $\pi : \Xi^{\mathbb{M}} \to \mathbb{M}$ is the map that forgets labels.

The following example is implicit in the literature, but is not usually named and does not have a consistent symbolic representation. We will use it enough that we must name it:

Example 5 (Metric thinning)**.** Let $\delta > 0$ be a tolerance parameter. The *δ-thinning* is the equivariant map $\theta_\delta : \mathbb{M} \to \mathbb{M}$ given by

$$\theta^\delta(\omega) = \{g \in \omega \mid d(g, \omega \setminus \{g\} > \delta\}.$$

[4] That is, a sum of Poisson-many independent random variables.

When θ^δ is applied to a point process, the result is always an δ-separated point process (but possibly empty).

We define θ^δ in the same way for marked point processes (that is, it simply ignores the marks).

Example 6 (Independent thinning)**.** Let Π be a point process. The *independent p-thinning* defined on its IID $[0,1]^\Pi$ is given by

$$\mathcal{I}_p([0,1]^\Pi) = \{g \in \Pi \mid \Pi_g \leq p\}.$$

One can show that independent p-thinning of the Poisson point process of intensity $t > 0$ yields the Poisson point process of intensity pt, as one would expect. See [cite] for further details.

Example 7 (Constant thickening)**.** Let $F \subset G$ be a finite set containing the identity $0 \in G$, and Π be a point process which is F*-separated* in the sense that $\Pi \cap \Pi f = \varnothing$ for all $f \in F \backslash \{0\}$. Then there is the associated thickening $\Theta^F(\Pi) = \Pi F$. It is intuitively obvious that $\operatorname{intensity} \Theta^F(\Pi) = |F| \operatorname{intensity} \Pi$. This can be formally established as follows: let $U \subseteq G$ be of unit volume. Then

$$\begin{aligned}
\operatorname{intensity} \Theta^F(\Pi) &= \mathbb{E}[|U \cap \Pi F|] && \text{By definition} \\
&= \sum_{f \in F} \mathbb{E}[|U \cap \Pi f|] && \text{By } F\text{-separation} \\
&= \sum_{f \in F} \mathbb{E}[\left|U f^{-1} \cap \Pi\right|] \\
&= \sum_{f \in F} \mathbb{E}[|U \cap \Pi|] && \text{By } \textit{unimodularity} \\
&= |F| \operatorname{intensity} \Pi.
\end{aligned}$$

This is the first real appearance of our unimodularity assumption. As is often the case, it is the formal reasoning behind obvious-looking statements.

In particular, we can demonstrate that $\operatorname{intensity} \Theta^F(\Pi) = |F| \operatorname{intensity} \Pi$ is *not* automatically true without unimodularity. For this, let Π denote the unit intensity Poisson point process on G, and $F = \{0, f\}$ where $f \in G$ is chosen such that $\lambda(Uf^{-1}) < 1$. Then $|Uf^{-1} \cap \Pi|$ is Poisson distributed with parameter $\lambda(Uf^{-1})$, and so by the above calculation $\operatorname{intensity} \Theta^F(\Pi) < 2 \cdot \operatorname{intensity} \Pi$.

Monotone maps have been investigated in the specific case of the Poisson point process on $\mathbb{R}^n$. We note the following interesting theorems:

Theorem 9 (Holroyd, Peres, Soo [HLS11])**.** Let $\lambda > \lambda' > 0$ be intensities. Then the Poisson point process on $\mathbb{R}^n$ of intensity λ can be thinned to the Poisson point process of intensity λ'. That is, there exists an equivariant and deterministic map $\theta : (\mathbb{M}(\mathbb{R}), \mathscr{P}_\lambda) \to (\mathbb{M}(\mathbb{R}), \mathscr{P}_{\lambda'})$.

Theorem 10 (Gurel-Gurevich and Peled [GGP13])**.** Let $\lambda > \lambda' > 0$ be intensities. Then the Poisson point process on $\mathbb{R}^n$ of intensity λ' cannot be thinned to the Poisson point process of intensity λ. That is, there is no equivariant and deterministic map $\Theta : (\mathbb{M}(\mathbb{R}), \mathscr{P}'_\lambda) \to (\mathbb{M}(\mathbb{R}), \mathscr{P}_\lambda)$.

Remark 8. The following two examples should be explored over nonamenable groups G.

We stress in the above theorems the *deterministic* nature of the maps. If one is allowed additional randomness (ie. one asks for a factor of IID map), then both constructions are trivially possible.

We note the following fact, which we will use (and prove) later after developing some notation.

Example 8. If Π is any point process, then its IID factors onto the Poisson (in fact, onto the IID Poisson).

Definition 9. A *factor* Ξ*-marking* of a point process is a G-equivariant map $\mathscr{C} : \mathbb{M} \to \Xi^{\mathbb{M}}$ such that the underlying subset in G of $\mathscr{C}(\omega)$ is ω. That is, $\mathscr{C}$ is a rule that assigns a mark from Ξ to each point of ω in some deterministic way. Again, if $\mathscr{C}$ is only defined μ almost everywhere then we will call it a μ *factor* Ξ*-marking*.

Example 9. Let $\theta : \mathbb{M} \to \mathbb{M}$ be a thinning. Then the associated 2-colouring is $\mathscr{C}_\theta : \mathbb{M} \to \{0,1\}^{\mathbb{M}}$ given by

$$\mathscr{C}_\theta(\omega) = \{(g, \mathbb{1}_{g \in \theta(\omega)} \in G \times \{0,1\} \mid g \in \omega\}.$$

We will see that all markings are built out of thinnings in a similar way.

Remark 9. There is a difference between the *thinning map* θ and the resulting *thinned process* $\theta_*(\mu)$ that can be a source for confusion. Passing to the thinned process (in principle) can lose information about μ.

For example, let Π denote a Poisson point process on G and Υ an independent random shift of a lattice $\Gamma < G$. Define the following thinning $\theta : \mathbb{M} \to \mathbb{M}$ by

$$\theta(\omega) = \{g \in \omega \mid g\Gamma \subseteq \omega\}.$$

Then $\theta(\Pi \cup \Upsilon) = \Upsilon$, and so the thinning completely loses the Poisson point process.

Definition 10. Let $\Phi : \mathbb{M} \to \mathbb{M}$ be a factor map. We think of its input ω as being red, its output $\Phi(\omega)$ as being blue, and their overlap $\omega \cap \Phi(\omega)$ as being purple.

For $g \in \omega$, let $\texttt{Colour}(g) \in \{\text{Red, Blue, Purple}\}$ be

$$\texttt{Colour}(g) = \begin{cases} \text{Red} & \text{If } g \in \omega \setminus \Phi(\omega), \\ \text{Blue} & \text{If } g \in \Phi(\omega) \setminus \omega, \\ \text{Purple} & \text{If } g \in \omega \cap \Phi(\omega). \end{cases}$$

Now define $\Theta^\Phi : \mathbb{M} \to \{\text{Red, Blue, Purple}\}^{\mathbb{M}}$ to be the following *input/output thickening* of Φ defined by

$$\Theta^\Phi(\omega) = \{(g, \texttt{Colour}(g)) \in G \times \text{Red, Blue, Purple}\} \mid g \in \omega\}.$$

Let $\pi : \{\text{Red, Blue, Purple}\}^{\mathbb{M}} \to \mathbb{M}$ be the projection map that deletes red points and then forgets colours, that is,

$$\pi(\omega) = \{g \in \omega \mid \omega_g \in \{\text{Blue, Purple}\}\}.$$

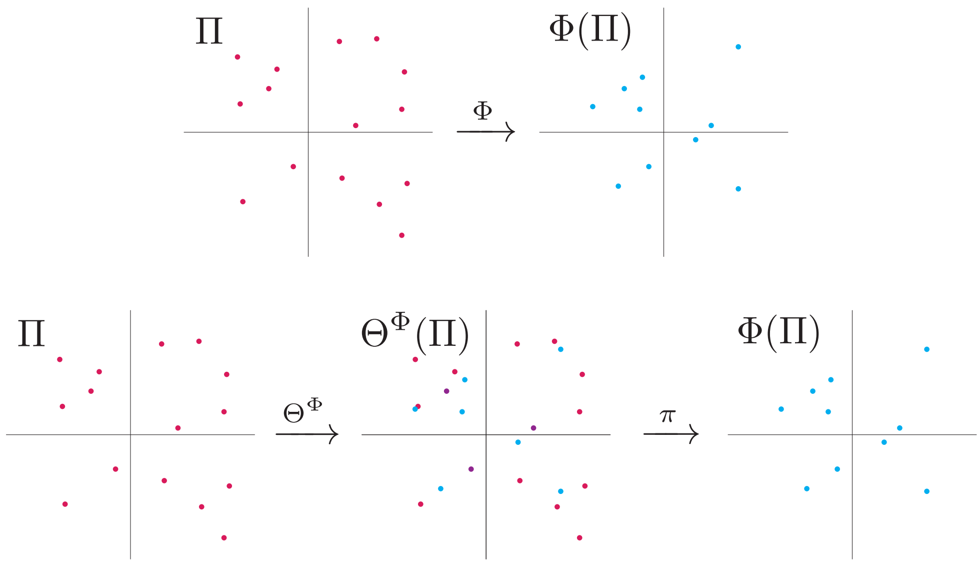

Figure 1.2: This is how you should picture the input/output thickening of a factor map.

Remark 10. Observe that $\Phi = \pi \circ \Theta^\Phi$ – that is, an *arbitrary* factor map decomposes as the composition of a thinning and a thickening. In this way we can often reduce the study of arbitrary factors to that of *monotone* factors.

Definition 11. The *space of graphs in G* is

$$\text{Graph}(G) = \{(V, E) \in \mathbb{M}(G) \times \mathbb{M}(G \times G) \mid E \subseteq V \times V\}.$$

This is a Borel G-space (with the diagonal action).

A *factor graph* is a measurable and G-equivariant map $\Phi : \mathbb{M}(G) \to \text{Graph}(G)$ with the property that the vertex set of $\Phi(\omega)$ is ω.

If a factor graph is connected, then we will refer to it as a *graphing*.

Remark 11. The elements of $\text{Graph}(G)$ are technically directed graphs, possibly with loops, and without multiple edges between the same pair of vertices. It's possible to define (in a Borel way) whatever space of graphs one desires (coloured, undirected, etc.) by taking appropriate subsets of products of configuration spaces.

Remark 12. One might prefer to call factor graphs as above *monotone* factor graphs. Our terminology follows that of [probabilists]. We have not found a use for the less restrictive factor graph concept.

Example 10. The *distance-R* factor graph is the map $\mathscr{D}_R : \mathbb{M} \to \text{Graph}(G)$ given by

$$\mathscr{D}_R(\omega) = \{(g, h) \in \omega \times \omega \mid d(g, h) \leq R\}.$$

The connectivity properties of this graph fall under the purview of continuum percolation theory.

1.2 The rerooting equivalence relation and groupoid

We now introduce a pair of algebraic objects that capture factors of a point process. For exposition's sake, we will first discuss unmarked point processes on a group G.

Definition 12. The *space of rooted configurations on G* is

$$\mathbb{M}_0(G) = \{\omega \in \mathbb{M}(G) \mid 0 \in \omega\}.$$

If G is understood, then we will drop it from the notation for clarity.

The *rerooting equivalence relation* on $\mathbb{M}_0$ is the orbit equivalence relation of $G \curvearrowright \mathbb{M}$ restricted to $\mathbb{M}_0$. Explicitly:

$$\mathcal{R} = \{(\omega, g^{-1}\omega) \in \mathbb{M}_0 \times \mathbb{M}_0 \mid g \in \omega\}.$$

This defines a countable Borel equivalence relation structure on $\mathbb{M}_0$. It is degenerate whenever $\omega \in \mathbb{M}_0$ exhibits symmetries: for instance, the equivalence class of $\mathbb{Z}$ viewed as an element of $\mathbb{M}_0(\mathbb{R})$ is a singleton. We are usually interested in essentially free actions, where such difficulties will not occur. Nevertheless, we do care about lattice shift point processes and so we will introduce a groupoid structure that keeps track of symmetries.

The *space of birooted configurations* is

$$\overrightarrow{\mathbb{M}_0} = \{(\omega, g) \in \mathbb{M}_0 \times G \mid g \in \omega\}.$$

We visualise an element $(\omega, g) \in \overrightarrow{\mathbb{M}_0}$ as the rooted configuration $\omega \in \mathbb{M}_0$ with an arrow pointing to $g \in \omega$ from the root (ie, the identity element of G).

The above spaces form a *groupoid* $(\mathbb{M}_0, \overrightarrow{\mathbb{M}_0})$ which we will refer to as the *rerooting groupoid*. Its unit space is $\mathbb{M}_0$ and its arrow space is $\overrightarrow{\mathbb{M}_0}$. We can identify $\mathbb{M}_0$ with $\mathbb{M}_0 \times \{0\} \subset \overrightarrow{\mathbb{M}_0}$.

The multiplication structure is as follows: we declare a pair of birooted configurations $(\omega, g), (\omega', h)$ in $\overrightarrow{\mathbb{M}_0}$ to be *composable* if $\omega' = g^{-1}\omega$, in which case

$$(\omega, g) \cdot (\omega', h) := (\omega, gh).$$

Note that if $\Gamma < G$ is a discrete subgroup (so $\Gamma \in \mathbb{M}_0(G)$), then the above multiplication is just the usual one.

The *source map* $s : \overrightarrow{\mathbb{M}_0} \to \mathbb{M}_0$ and *target map* $t : \overrightarrow{\mathbb{M}_0} \to \mathbb{M}_0$ are

$$s(\omega, g) = \omega, \text{ and } t(\omega, g) = g^{-1}\omega.$$

Note that the rerooting groupoid is *discrete* in the sense that $s^{-1}(\omega)$ is at most countable for all $\omega \in \mathbb{M}_0$.

Remark 13. Let $\mathbb{M}_0^{\text{aper}}$ denote the set of rooted configurations ω that are *aperiodic* in the sense that $\text{stab}_G(\omega) = \{e\}$. Then the groupoid generated by $\mathbb{M}_0^{\text{aper}}$ in $\mathbb{M}_0$ is *principal*.

Definition 13. If Ξ is a space of marks, then the *space of Ξ-marked rooted configurations* is

$$\Xi^{\mathbb{M}_0} = \{\omega \in \Xi^{\mathbb{M}} \mid \exists \xi \in \Xi \text{ such that } (0, \xi) \in \omega\}.$$

The *Ξ-marked rerooting groupoid* is defined as previously, with $\Xi^{\mathbb{M}_0}$ taking the place of $\mathbb{M}_0$.

1.2.1 Borel correspondences between the groupoid and factors

Suppose $\theta : \mathbb{M} \to \mathbb{M}$ is an equivariant and measurable thinning. Then we can associate to it a subset of the rerooting groupoid, namely

$$A_\theta = \{\omega \in \mathbb{M} \mid 0 \in \theta(\omega)\}.$$

This association has an inverse: given a subset $A \subseteq \mathbb{M}_0$, we can define a thinning $\theta^A : \mathbb{M} \to \mathbb{M}$

$$\theta^A(\omega) = \{g \in \omega \mid g^{-1}\omega \in A\}.$$

Thus we see that *Borel subsets $A \subseteq \mathbb{M}_0$ of the rerooting groupoid correspond to Borel thinning maps $\theta : \mathbb{M} \to \mathbb{M}$.*

In the Ξ-marked case, one associates to a subset $A \subseteq \Xi^{\mathbb{M}_0}$ a thinning $\theta^A : \Xi^{\mathbb{M}} \to \Xi^{\mathbb{M}}$.

In a similar way, we can see that if $P : \mathbb{M}_0 \to [d]$ is a Borel partition of $\mathbb{M}_0$ into d classes, then there is an associated factor $[d]$-colouring $\mathscr{C}^P : \mathbb{M} \to [d]^{\mathbb{M}}$ given by

$$\mathscr{C}^P(\omega) = \{(g, P(g^{-1}\omega) \in G \times [d] \mid g \in \omega\},$$

and given a factor $[d]$-colouring $\mathscr{C} : \mathbb{M} \to [d]^{\mathbb{M}}$ one associates the partition $P^{\mathscr{C}} : \mathbb{M}_0 \to [d]$ given by

$$P(\omega) = c, \text{ where } c \text{ is the unique element of } [d] \text{ such that } (0, c) \in \mathscr{C}(\omega).$$

Again, these associations are mutual inverses. Now suppose that $\mathscr{G} : \mathbb{M} \to \mathrm{Graph}(G)$ is an equivariant and measurable factor graph. Then we can associate to it a subset of the rerooting groupoid's arrow space, namely

$$\mathscr{A}_{\mathscr{G}} = \{(\omega, g) \in \overrightarrow{\mathbb{M}_0} \mid (\omega, g^{-1}\omega) \in \mathscr{G}(\omega)\}.$$

In the other direction, we associate to a subset $\mathscr{A} \subseteq \overrightarrow{\mathbb{M}_0}$ the factor graph $\mathscr{G}^{\mathscr{A}} : \mathbb{M} \to \mathrm{Graph}(G)$

$$\mathscr{G}^{\mathscr{A}}(\omega) = \{(g, h) \in \omega \times \omega \mid (g^{-1}\omega, g^{-1}h) \in \mathscr{A}\}.$$

Thus we see that *Borel subsets $\mathscr{A} \subseteq \overrightarrow{\mathbb{M}_0}$ of the rerooting groupoid's arrow space correspond to Borel factor graphs $\mathscr{G} : \mathbb{M} \to \mathrm{Graph}(G)$.*

Note also that the factor graph $\mathscr{G}$ is *connected* for every input ω if and only if the corresponding subset $\mathscr{A}_{\mathscr{G}} \subseteq \overrightarrow{\mathbb{M}_0}$ generates the rerooting groupoid.

If μ is a point process, then the correspondence still works in one direction: namely, we can associate subsets $A \subset \mathbb{M}_0$ (or $\mathscr{A} \subseteq \overrightarrow{\mathbb{M}_0}$) to μ thinnings $\theta^A : (\mathbb{M}, \mu) \to \mathbb{M}$ (or μ factor graphs $\mathscr{G}_{\mathscr{A}} : (\mathbb{M}, \mu) \to \mathbb{M}$ respectively). It turns out that we can make sense of these correspondences in the measured category, but we will require some theory developed in the next section.

1.3 The Palm measure

We will now associate to a (finite intensity) point process μ a probability measure μ_0 defined on the rerooting groupoid $\mathbb{M}_0$. When the ambient space is unimodular, this will turn the rerooting groupoid into a *probability measure preserving (pmp) discrete groupoid.*

Informally, the Palm measure of a point process Π is the process conditioned to contain the root. A priori this makes no sense (the subset $\mathbb{M}_0$ has probability zero), but there is an obvious way one could interpret the statement: condition on the process to contain a point in an ε ball about the root, and take the limit as ε goes to zero. See Theorem 13.3.IV of [DVJ07] and Section 9.3 of [LP18] for further details.

We will instead take the following *relative rate* as our basic definition:

Definition 14. Let Π be a point process of finite intensity with law μ. Its (normalised) *Palm measure* is the probability measure μ_0 defined on $\mathbb{M}_0$ given by

$$\mu_0(A) := \frac{\text{intensity}(\theta^A(\Pi))}{\text{intensity}(\Pi)},$$

where θ^A is the thinning associated to $A \subseteq \mathbb{M}_0$.

More explicitly,

$$\mu_0(A) := \frac{1}{\text{intensity}(\mu)} \mathbb{E}_\mu \left[\#\{g \in U \mid g^{-1}\omega \in A\} \right],$$

where $U \subseteq G$ is of unit volume.

We also define the Palm measure of a Ξ-marked point process similarly, with $\Xi^{\mathbb{M}_0}$ taking the place of $\mathbb{M}_0$.

A *Palm version* of Π is any random variable Π_0 with law μ_0. That is, if for all $B \subseteq \mathbb{M}_0$ we have

$$\mathbb{P}[\Pi_0 \in B] = \mu_0(B).$$

We now give examples of how to express the Palm measure of various examples of point processes. We will require all of these examples, so give proofs with one exception (whose proof requires more point process theory than we wish to give, and whose statement is simple and intuitive anyway).

Example 11 (Forgetting labels)**.** When talking about the Palm measure for a Ξ-marked point process, it is important in the above to choose the correct thinning. Recall from Remark 7 that for a subset $A \subseteq \Xi^{\mathbb{M}_0}$ one can discuss *two* possible kinds of thinnings, namely

$$\theta^A : \Xi^{\mathbb{M}} \to \Xi^{\mathbb{M}} \text{ or } \pi \circ \theta^A : \Xi^{\mathbb{M}} \to \mathbb{M},$$

where $\pi : \Xi^{\mathbb{M}} \to \mathbb{M}$ is the map that forgets labels.

It is the *former* kind of thinning one should take.

Note that if Π is a Ξ-marked point process, then its intensity remains the same if you forget the marks, that is, intensity Π = intensity $\pi(\Pi)$. More generally, the operation of taking the Palm version *commutes* with forgetting labels. That is, $\pi(\Pi_0) = (\pi(\Pi))_0$. To see this, let $B \subseteq \mathbb{M}_0$, and observe

$$\begin{aligned}
\mathbb{P}[\pi(\Pi_0) \in B] &= \mathbb{P}[\Pi_0 \in \pi^{-1}(B)] \\
&= \frac{\text{intensity } \theta^{\pi^{-1}(B)}(\Pi)}{\text{intensity } \Pi} = \frac{\text{intensity } \pi(\theta^{\pi^{-1}(B)}(\Pi))}{\text{intensity } \pi(\Pi)} = \frac{\text{intensity } \theta^B(\pi(\Pi))}{\text{intensity } \pi(\Pi)} \\
&= \mathbb{P}[\pi(\Pi)_0 \in B],
\end{aligned}$$

where we simply follow our nose.

Example 12 (Lattice actions)**.** If $\Gamma < G$ is a lattice, then the Palm measure of the associated lattice shift is just δ_Γ – that is, the atomic measure on $\Gamma \in \mathbb{M}_0(G)$. More generally, if $\Gamma \curvearrowright (X, \mu)$ is a pmp action, then the Palm measure of the associated induced X-marked point process is its *symbolic dynamics*. That is, the map $\Sigma : (X, \mu) \to X^{\mathbb{M}}$ given by

$$\Sigma(x) = \{(\gamma, \gamma^{-1} \cdot x) \in G \times X \mid \gamma \in \Gamma\}.$$

pushes forward μ to the Palm measure. In words, you sample a μ-random point $x \in X$ and track its orbit under Γ (the inverse is an artefact of our left bias).

Example 13 (Mecke-Slivnyak Theorem)**.** If Π is a Poisson point process, then its Palm measure has the same law as $\Pi \cup \{0\}$, where $0 \in G$ is the identity.

The proof of the above fact can be found in Section 9.2 of [LP18]. In fact, this is a *characterisation* of the Poisson point process: if the Palm measure of μ is obtained by simply adding the root[5], then μ is the Poisson point process (of some intensity).

As a consequence, the Palm measure of the IID Poisson is the IID of the Palm measure of the Poisson itself.

Example 14 (Thinnings)**.** The Palm version $\theta(\Pi)_0$ of a thinning $\theta = \theta^A$ of Π (determined by a subset $A \subseteq \mathbb{M}_0$) is described in terms of its Palm version Π_0 as a conditional probability as follows:

$$\mathbb{P}[\theta(\Pi)_0 \in B] = \mathbb{P}[\theta(\Pi_0) \in B \mid \Pi_0 \in A]$$

for any $B \subseteq \mathbb{M}_0$.

To see this, first one should work from the definitions to show that $\theta^B(\theta^A(\Pi)) = \theta^{A \cap (\theta^A)^{-1}(B)}$. Therefore

$$\begin{aligned}
\mathbb{P}[(\theta(\Pi))_0 \in B] &= \frac{\text{intensity}\, \theta^B(\theta^A(\Pi))}{\text{intensity}\, \theta^A(\Pi)} \\
&= \frac{\text{intensity}\, \theta^{A \cap (\theta^A)^{-1}(B)}(\Pi)}{\text{intensity}\, \Pi} \Bigg/ \frac{\text{intensity}\, \theta^A(\Pi)}{\text{intensity}\, \Pi} \qquad \text{By the observation} \\
&= \frac{\mathbb{P}[\Pi_0 \in A \cap (\theta^A)^{-1}(B)]}{\mathbb{P}[\Pi_0 \in A]} \\
&= \frac{\mathbb{P}[\{\theta(\Pi_0) \in B\} \cap \{\Pi_0 \in A\}]}{\mathbb{P}[\Pi_0 \in A]} \\
&= \frac{\mathbb{P}[\{\theta(\Pi_0) \in B\} \cap \{\Pi_0 \in A\}]}{\mathbb{P}[\Pi_0 \in A]},
\end{aligned}$$

which is exactly the definition of the desired conditional probability.

Example 15. The Palm version $\mathscr{C}(\Pi)_0$ of a 2-colouring $\mathscr{C} : \mathbb{M} \to \{0,1\}^{\mathbb{M}}$ determined by a subset $A \subseteq \mathbb{M}_0$ can be described as follows. Let us write $\Pi = \Pi^0 \sqcup \Pi^1$, where Π^i denotes the subset of Π's points coloured i. We also let $\mathbb{M}_0^i$ denote the collection of rooted configurations where the root is coloured i.

One can readily compute that $\pi(\theta^{B \cap \mathbb{M}_0^i}(\mathscr{C}(\Pi))) = \theta^{\mathscr{C}^{-1}(B \cap \mathbb{M}_0^i)}(\Pi)$, where $\pi : \{0,1\}^{\mathbb{M}} \to \mathbb{M}$ is the map that forgets labels.

[5] More formally, consider the map $F : \mathbb{M} \to \mathbb{M}_0$ given by $F(\omega) = \omega \cup \{0\}$, by "adding the root" we mean the Palm measure μ_0 is the pushforward $F_* \mu$.

For $B \subseteq \{0,1\}^{\mathbb{M}_0}$ we have

$$\begin{aligned}
\mathbb{P}[\mathscr{C}(\Pi)_0 \in B] &= \mathbb{P}[\mathscr{C}(\Pi)_0 \in B \cap \mathbb{M}_0^0] + \mathbb{P}[\mathscr{C}(\Pi)_0 \in B \cap \mathbb{M}_0^1] \\
&= \frac{\text{intensity}\, \theta^{B\cap\mathbb{M}_0^0}(\mathscr{C}(\Pi))}{\text{intensity}\, \mathscr{C}(\Pi)} + \frac{\text{intensity}\, \theta^{B\cap\mathbb{M}_0^1}(\mathscr{C}(\Pi))}{\text{intensity}\, \mathscr{C}(\Pi)} \\
&= \frac{\text{intensity}\, \pi(\theta^{B\cap\mathbb{M}_0^0}(\Pi))}{\text{intensity}\, \Pi} + \frac{\text{intensity}\, \pi(\theta^{B\cap\mathbb{M}_0^1}(\Pi))}{\text{intensity}\, \Pi} \\
&= \frac{\text{intensity}\, \theta^{\mathscr{C}^{-1}(B\cap\mathbb{M}_0^0)}(\Pi)}{\text{intensity}\, \Pi} + \frac{\text{intensity}\, \theta^{\mathscr{C}^{-1}(B\cap\mathbb{M}_0^1)}(\Pi)}{\text{intensity}\, \Pi} \\
&= \mathbb{P}[\Pi_0 \in \mathscr{C}^{-1}(B \cap \mathbb{M}_0^0)] + \mathbb{P}[\Pi_0 \in \mathscr{C}^{-1}(B \cap \mathbb{M}_0^1)] \\
&= \mathbb{P}[\{\mathscr{C}(\Pi_0) \in B\} \cap \{(0,0) \in \Pi_0\}] + \mathbb{P}[\{\mathscr{C}(\Pi_0) \in B\} \cap \{(0,1) \in \Pi_0\}] \\
&= \mathbb{P}[\mathscr{C}(\Pi_0) \in B \mid (0,0) \in \Pi_0]\mathbb{P}[(0,0) \in \Pi_0] \\
&\quad + \mathbb{P}[\mathscr{C}(\Pi_0) \in B \mid (0,1) \in \Pi_0]\mathbb{P}[(0,1) \in \Pi_0].
\end{aligned}$$

That is to say, the Palm version of $\mathscr{C}(\Pi)$ is the $\mathscr{C}(\Pi_0)$ where you choose Π_0 proportional to the frequency of the colour at the root.

Example 16. Let $\Theta = \Theta^F$ be a constant thickening determined by $F \subset G$, as described in Example 7. If Π is an F-separated process, then the Palm version $\Theta(\Pi)_0$ of the thickening $\Theta(\Pi)$ is as follows: sample from Π_0, and independently choose to root $\Theta(\Pi_0)$ at a uniformly chosen element X of F. That is, $\Theta(\Pi)_0 \overset{d}{=} X^{-1}\Theta(\Pi_0)$.

To see this, we compute[6] as follows:

$$\begin{aligned}
\mathbb{P}[\Theta(\Pi)_0 \in B] &= \frac{1}{\text{intensity}\, \Theta(\Pi)} \mathbb{E}[\#\{g \in U \cap \Pi F \mid g^{-1}\Theta(\Pi) \in B\}] && \text{By definition} \\
&= \frac{1}{|F|}\frac{1}{\text{intensity}\, \mu} \sum_{f\in F} \mathbb{E}[\#\{g \in U \cap \Pi f \mid g^{-1}\Theta(\Pi) \in B\}] && \text{By Example 7} \\
&= \frac{1}{|F|}\frac{1}{\text{intensity}\, \mu} \sum_{f\in F} \mathbb{E}[\#\{g \in Uf^{-1} \cap \Pi \mid g^{-1}\Pi \in \Theta^{-1}(B)\}] && \text{By equivariance} \\
&= \frac{1}{|F|}\frac{1}{\text{intensity}\, \mu} \sum_{f\in F} \mathbb{E}[\#\{g \in U \cap \Pi \mid g^{-1}\Pi \in \Theta^{-1}(B)\}] && \text{By unimodularity} \\
&= \frac{1}{|F|} \sum_{f\in F} \mathbb{P}[\Pi_0 \in \Theta^{-1}(B)] && \text{By definition} \\
&= \frac{1}{|F|} \sum_{f\in F} \mathbb{P}[\Theta(\Pi_0) \in B] \\
&= \mathbb{P}[X^{-1}\Theta(\Pi_0) \in B].
\end{aligned}$$

The Palm measure has an associated integral equation. One writes $(\lambda \otimes \mu_0)(U \times A) = \int_G \mathbb{E}_0[\mathbb{1}_{U\times A}]d\lambda(x)$, and then invokes the usual voodoo to extend a statement about measurable sets to one about measurable functions. The resulting theorem is:

[6] When we define the Palm measure of a set $B \subseteq \mathbb{M}_0$, we usually write "$g \in U$" rather than "$g \in U \cap \Pi$", as the latter condition $g^{-1}\Pi \in B$ already implies $g \in \Pi$. For this computation it is better to really spell it out though.

Theorem 11 (Campbell-Little-Mecke-Matthes)**.** Let μ be a finite intensity point process on G with Palm measure μ_0. Write $\mathbb{E}$ and $\mathbb{E}_0$ for the associated integral operators.

If $f : G \times \mathbb{M}_0 \to \mathbb{R}_{\geq 0}$ is a measurable function (*not* necessarily invariant in any way), then

$$\mathbb{E}\left[\sum_{x \in \omega} f(x, x^{-1}\omega)\right] = \text{intensity}(\mu)\mathbb{E}_0\left[\int_G f(x,\omega)d\lambda(x)\right].$$

Note that summing against ω is the same as integrating G against ω viewed as a locally finite measure on G.

Remark 14. There is a map $\mathscr{V} : [0,1] \times \mathbb{M}_0 \to \mathbb{M}$ with the property that if μ is *any* point process with Palm measure μ_0, then $V_*(\texttt{Leb} \otimes \mu_0) = \mu$. This is a consequence of the *Voronoi inversion formula*, see Section 9.4 of [LP18]. Thus if ν is a point process with $\nu_0 = \mu_0$, then $\nu = \mu$, that is, the Palm measure *determines* the point process.

1.3.1 Unimodularity and the Mass Transport Principle

The source and range maps $s, t : \overrightarrow{\mathbb{M}_0} \to \mathbb{M}$ induce a pair of measures on $\overrightarrow{\mathbb{M}_0}$ defined by

$$\overrightarrow{\mu_0}^s(\mathscr{G}) = \int_{\mathbb{M}_0} \left|s^{-1}(\omega) \cap \mathscr{G}(\omega)\right| d\mu_0(\omega), \text{ and } \overrightarrow{\mu_0}^t(\mathscr{G}(\omega)) = \int_{\mathbb{M}_0} \left|t^{-1}(\omega) \cap \mathscr{G}\right| d\mu_0(\omega).$$

In our factor graph interpretation, this corresponds to the expected indegree / outdegree of $\mathscr{G}$ respectively, where we view $\mathscr{G}$ as a *directed* graph. To see this, recall that for a rooted configuration $\omega \in \mathbb{M}_0$,

$$s^{-1}(\omega) = \{(\omega, g) \in \mathbb{M}_0 \times G \mid g \in \omega\} \text{ and } t^{-1}(\omega) = \{(g^{-1}\omega, g^{-1}) \in \mathbb{M}_0 \times G \mid g \in \omega\},$$

and that there is an edge from 0 to g in $\mathscr{G}(\omega)$ exactly when $(\omega, g) \in \mathscr{G}$, and an edge from g to 0 exactly when $(g^{-1}\omega, g^{-1}) \in \mathscr{G}$. Thus

$$\overrightarrow{\deg}_0(\mathscr{G}(\omega)) = \left|s^{-1}(\omega) \cap \mathscr{G}(\omega)\right| \text{ and } \overleftarrow{\deg}_0(\mathscr{G}(\omega)) = \left|t^{-1}(\omega) \cap \mathscr{G}(\omega)\right|.$$

Remark 15. We have had to adapt notation to suit our purposes. Usually a groupoid would be denoted by a letter like $\mathcal{G}$, and that is the set of arrows. Then its units would be denoted $\mathcal{G}_0$. We have tried to match this up with the necessary notation from point process theory as closely as possible.

We choose to denote outdegree by an expression like $\overrightarrow{\deg}_0(\mathscr{G}(\omega))$ instead of $\deg^+_{\mathscr{G}(\omega)}(0)$ as the arrows are more evocative, and the subscript notation becomes very small (as in, for instance, $\deg^+_{\mathscr{G}(\Pi_0)}(0)$.

Proposition 4. If G is *unimodular*, then $\overrightarrow{\mu_0}^s = \overrightarrow{\mu_0}^t$. That is, $(\overrightarrow{\mathbb{M}_0}, \overrightarrow{\mu_0})$ forms a discrete pmp groupoid.

Equivalently, if Π_0 is the Palm version of any point process Π on G, then

$$\mathbb{E}\left[\overrightarrow{\deg}_0(\mathscr{G}(\Pi_0))\right] = \mathbb{E}\left[\overleftarrow{\deg}_0(\mathscr{G}(\Pi_0))\right].$$

We will denote by $\overrightarrow{\mu_0}$ this common measure $\overrightarrow{\mu_0}^s = \overrightarrow{\mu_0}^t$.

Definition 15. The *Palm groupoid* of a point process Π with law μ is $(\overrightarrow{\mathbb{M}_0}, \overrightarrow{\mu_0})$. If Π is free, then this groupoid is principal, and thus we refer to Π's *Palm equivalence relation* $(\mathbb{M}_0, \mathcal{R}, \mu_0)$.

To the author's knowledge, the only direct references in the literature to the existence of this equivalence relation can be found in a paper of Avni [Avn05] (Example 2.2) and a paper of Bowen [Bow18] (Questions and comments, item 1).

This makes a certain amount of sense: the kind of people interested in point processes are typically repulsed by cbers, and vice-versa.

There are also implicit references: see [DVJ07], [MI17], [BHM18].

Definition 16. Let Π be a point process and $\mathscr{G}$ an *undirected* factor graph of Π. Its *edge density* is $\mathbb{E}[\deg_0(\mathscr{G}(\Pi_0))$, where Π_0 is the Palm version of Π.

By the above proposition, if $\mathscr{G}'$ is any *orientation* of $\mathscr{G}$, then the edge density can be expressed as

$$\mathbb{E}[\deg_0(\mathscr{G}(\Pi_0)) = 2\mathbb{E}\left[\overrightarrow{\deg_0}(\mathscr{G}'(\Pi_0))\right].$$

Speaking properly then, we should be talking of *directed* factor graphs, but for this reason we will often think of the factor graphs as undirected.

Proof of Proposition 4.

$$\begin{aligned}
\overrightarrow{\mu_0}^s(\mathscr{G}) &= \mathbb{E}_{\mu_0}\left[\sum_{g\in\omega} \mathbb{1}_{(\omega,g)\in\mathscr{G}}\right] && \text{by definition} \\
&= \mathbb{E}_{\mu_0}\left[\int_G \mathbb{1}_{x\in U} \sum_{g\in\omega} \mathbb{1}_{(\omega,g)\in\mathscr{G}} d\lambda(x)\right] && \text{For any } U\subseteq G \text{ of unit volume} \\
&= \frac{1}{\text{intensity }\mu}\mathbb{E}_{\mu}\left[\sum_{x\in\omega} \mathbb{1}_{x\in U} \sum_{g\in x^{-1}\omega} \mathbb{1}_{(x^{-1}\omega,g)\in\mathscr{G}}\right] && \text{By the CLLM} \\
&= \frac{1}{\text{intensity }\mu}\mathbb{E}_{\mu}\left[\sum_{h\in\omega} \sum_{hg^{-1}\in\omega} \mathbb{1}_{hg^{-1}\in U} \mathbb{1}_{(gh^{-1}\omega,g)\in\mathscr{G}}\right] && \text{Fubini and variable change } h = xg \\
&= \mathbb{E}_{\mu_0}\left[\int_G \sum_{g\in\omega} \mathbb{1}_{h^{-1}g\in U} \mathbb{1}_{(g\omega,g)\in\mathscr{G}} d\lambda(h)\right] && \text{By the CLLM} \\
&= \mathbb{E}_{\mu_0}\left[\sum_{g\in\omega} \underbrace{\left(\int_G \mathbb{1}_{h^{-1}g\in U} d\lambda(h)\right)}_{=\lambda((Ug)^{-1})} \mathbb{1}_{(g\omega,g)\in\mathscr{G}}\right] && \text{Fubini} \\
&= \mathbb{E}_{\mu_0}\left[\sum_{g\in\omega} \mathbb{1}_{(g\omega,g)\in\mathscr{G}}\right] && \text{By unimodularity} \\
&= \overrightarrow{\mu_0}^t(\mathscr{G}).
\end{aligned}$$

□

This directly implies by the usual voodoo for extending a statement about equality of measures to equality of integrals to:

Theorem 12 (The Mass Transport Principle)**.** Let μ be a point process on a unimodular group. Suppose $T : G \times G \times \mathbb{M} \to \mathbb{R}_{\geq 0}$ is a measurable function which is *diagonally invariant* in the sense that $T(gx, gy; g\omega) = T(x, y; \omega)$ for all $g \in G$. Then

$$\mathbb{E}_{\mu_0}\left[\sum_{x\in\omega} T(x, 0; \omega)\right] = \mathbb{E}_{\mu_0}\left[\sum_{y\in\omega} T(0, y; \omega)\right].$$

We view $T(x, y; \omega)$ as representing an amount of *mass* sent from x to y when the configuration is ω. Thus the integrand on the lefthand side represents the total mass received from the root, and similarly the integrand on the righthand side represents the total mass sent from the root.

The mass transport principle immediately follows from Proposition 4, as it just represents the integral of the function $\omega \mapsto \sum_{x\in\omega} T(x, 0; \omega)$ with respect to $\overrightarrow{\mu_0}^t$ and $\overrightarrow{\mu_0}^s$.

Remark 16. One can use the CLLM formula to express $\overrightarrow{\mu_0}(\mathscr{G})$ without reference to the Palm measure. Let $U \subseteq G$ be of unit volume, and apply the formula to $f(x, \omega) = \mathbb{1}_{x\in U}\overrightarrow{\deg}_0(\mathscr{G}(\omega))$, resulting in

$$\overrightarrow{\mu_0}(\mathscr{G}) = \frac{1}{\text{intensity}\,\Pi}\mathbb{E}\left[\sum_{x\in\Pi} \mathbb{1}_{x\in U}\overrightarrow{\deg}_x(\mathscr{G}(\Pi))\right]$$

(note that by equivariance $\overrightarrow{\deg}_0(\mathscr{G}(x^{-1}\omega)) = \deg_x(\mathscr{G}(\omega))$).

Example 17 (General thickening)**.** Suppose one has for each configuration $\omega \in \mathbb{M}$ and each $g \in \omega$ a *finite* subset $F_\omega(g)$ satisfying the following properties:

Separation: If $g, h \in \omega$ are *distinct* then $F_\omega(g) \cap F_\omega(h) = \varnothing$, and

Equivariance: For all $\gamma \in G$, we have $F_{\gamma\omega}(\gamma\omega) = \gamma F_\omega(g)$.

Then one can define a thickening $\Theta : \mathbb{M} \to \mathbb{M}$ by

$$\Theta(\omega) = \bigsqcup_{g\in\omega} gF_\omega(g).$$

That is, each point $g \in \omega$ looks at the current configuration, and adds points $F_\omega(g)$ local to it according to some equivariant rule.

It stands to reason that if Π is a point process satisfying the above rules almost surely, then intensity $\Theta(\Pi) = \mathbb{E}|F_{\Pi_0}| \cdot$ intensity Π. Just as in Example 7 though, this will require unimodularity to prove.

It will be convenient to work not with Θ directly, but with the following $\mathbb{N}_0$-coloured version of it

$$\Theta^C(\omega) = \{(g, F_\omega(g)) \in G \times \mathbb{N} \mid g \in \omega\} \sqcup \Theta(\omega) \times \{0\}.$$

Here we are simply colouring each point according to how many points it added, and colouring these added points by 0. The point is that the resulting map Θ^C is *injective*. Note that intensity$(\Theta^C(\Pi))$ = intensity$(\Theta(\Pi))$. Let us write θ for the map that takes a $\mathbb{N}_0$-marked configuration and spits out its set of points with labels *not* equal to zero. This is a thinning.

Let us define the following *transport function* in order to apply the mass transport principle:

$$T(x,y;\Theta^C(\omega)) = \begin{cases} 1 & \text{If } x \in F_\omega(y), \\ 0 & \text{else.} \end{cases}$$

Observe that $\sum_{y\in\Theta^C(\omega)} T(0,y;\Theta^C(\omega))$ is identically one, as each point sends unit mass to exactly one point. On the other side,

$$\sum_{x\in\Theta^C(\omega)} T(x,0;\Theta^C(\omega)) = \begin{cases} |F_\omega(0)| & 0 \in \omega \\ 0 & \text{else.} \end{cases}$$

We have to compute the expectation of this function *with respect to* $\Theta^C(\Pi)_0$. For this random variable, the condition that "$0 \in \omega$" translates to saying that the root should be $\mathbb{N}$-coloured. On that event, $\Theta^C(\Pi)_0$ is distributed according to $\Theta^C(\Pi_0)$. Then the expected mass in is

$$\mathbb{E}\left[\left|F_{\theta(\Theta^C(\Pi)_0)}(0)\right|\mathbb{1}[0\in\theta(\Theta^C(\Pi)_0)]\right] = \mathbb{E}\left[\left|F_{\theta(\Theta^C(\Pi)_0)}(0)\right|\middle|0\in\theta(\Theta^C(\Pi)_0)\right]\mathbb{P}[0\in\theta(\Theta^C(\Pi)_0)]$$
$$= \mathbb{E}\left[\left|F_{\theta(\Theta^C(\Pi))_0}(0)\right|\right]\frac{\text{intensity}\,\Pi}{\text{intensity}(\Theta(\Pi)}$$
$$= \mathbb{E}\left[|F_{\Pi_0}(0)|\right]\frac{\text{intensity}\,\Pi}{\text{intensity}(\Theta(\Pi))}.$$

So by the mass transport principle, this final term is equal to one.

This has been notationally cumbersome in the hopes of being precise. There must be a better way to express this. For our next trick we will skip the surplus notation and hope it is clear enough to follow.

Now we determine the Palm version of $\Theta^C(\Pi)$, and hence of $\Theta(\Pi)$ itself (take the former and forget its labels).

Let X be uniform over $F_{\Pi_0}(0)$, conditional on Π_0. Then

$$\mathbb{P}(\Theta^C(\Pi)_0 \in B] = \sum_{k\geq 1}\mathbb{P}[X^{-1}\Theta^C(\Pi_0)\in B\Big||F_{\Pi_0}(0)|=k]\mathbb{P}[|F_{\Pi_0}(0)|=k].$$

That is, the Palm version of $\Theta^C(\Pi)$ can be found from taking a *size-biased* Palm version of Π, applying Θ^C, and then rooting at random.

This is also proved by mass transport. Define for $B \subseteq \mathbb{N}_0^{\mathbb{M}_0}$

$$T(x,y;\Theta^C(\omega)) = \begin{cases} 1 & \text{If } x \in F_\omega(y)\cap\theta^B(\Theta^C(\omega)) \\ 0 & \text{else.} \end{cases}$$

Observe that $\sum_{y\in\Theta^C(\omega)} T(0,y;\Theta^C(\omega))$ is $\mathbb{1}[0\in\theta^B(\Theta^C(\omega))]$, which has expectation exactly $\mathbb{P}[\Theta^C(\Pi)_0\in B]$. On the other side,

$$\sum_{x\in\Theta^C(\omega)} T(x,0;\Theta^C(\omega)) = \begin{cases} \#\{x\in F_\omega(0) \mid x^{-1}\Theta^C(\omega)\in B\} & 0\in\omega \\ 0 & \text{else.} \end{cases}$$

This has expectation

$$\begin{aligned}
&\mathbb{E}[\#\{x \in F_{\Pi_0}(0) \mid x^{-1}\Theta^C(\Pi_0) \in B]\\
&= \sum_{k\geq 1} \mathbb{E}[\#\{x \in F_{\Pi_0}(0) \mid x^{-1}\Theta^C(\Pi_0) \in B \big| |F_{\Pi_0}(0)| = k]\mathbb{P}[|F_{\Pi_0}(0)| = k]\\
&= \sum_{k\geq 1} \mathbb{P}[X^{-1}\Theta^C(\Pi_0) \in B \big| |F_{\Pi_0}| = k]\mathbb{P}[|F_{\Pi_0}(0)| = k].
\end{aligned}$$

1.3.2 Ergodicity and the factor correspondences in the measured category

Definition 17. A subset $A \subseteq \mathbb{M}$ of unrooted configurations is *shift-invariant* if for all $\omega \in A$ and $g \in G$, we have $g\omega \in A$.

A subset $A_0 \subseteq \mathbb{M}_0$ of rooted configurations is *rootshift invariant* if for all $\omega \in A_0$ and $g \in \omega$, we have $g^{-1}\omega \in A_0$.

Note that if $A \subseteq \mathbb{M}$ is shift-invariant, then $A_0 := A \cap \mathbb{M}_0$ is rootshift invariant, and if $A_0 \subseteq \mathbb{M}_0$ is rootshift-invariant, then $A := GA_0$ is shift invariant. More is true:

Proposition 5. Let μ be a point process with Palm measure μ_0.

1. If $A \subseteq \mathbb{M}_0$ is rootshift invariant, then $\mu_0(A) = \mu(GA)$.
2. If $A \subseteq \mathbb{M}$ is shift invariant, then $\mu_0(A \cap \mathbb{M}_0) = \mu(A)$.

That is, under the correspondence between rootshift invariant subsets of $\mathbb{M}_0$ and shift invariant subsets of $\mathbb{M}$, the measures μ_0 and μ coincide.

In particular, $G \curvearrowright (\mathbb{M}, \mu)$ is ergodic *if and only if* $(\mathbb{M}_0, \mathcal{R}, \mu_0)$ is ergodic.

Proof. We assume ergodicity and prove the statements about measures. The general case will follow.

First, suppose $G \curvearrowright (\mathbb{M}, \mu)$ is ergodic, and let $A \subseteq \mathbb{M}_0$ be rootshift invariant. Then for any $U \subseteq G$ of unit volume,

$$\begin{aligned}
\mu_0(A) &= \frac{1}{\text{intensity}\,\mu}\mathbb{E}_\mu\left[\#\{g \in U \mid g^{-1}\omega \in A\}\right] && \text{By definition}\\
&= \frac{1}{\text{intensity}\,\mu}\mathbb{E}_\mu\left[|\omega \cap U|\mathbb{1}_{\omega\in GA}\right] && \text{By rootshift invariance of } A\\
&= \mu(GA) && \text{By ergodicity.}
\end{aligned}$$

In particular, we see that $\mu_0(A)$ is zero or one, so the equivalence relation is ergodic.

Now suppose $(\mathbb{M}_0, \mathcal{R}, \mu)$ is ergodic, and let $A \subseteq \mathbb{M}$ be shift invariant.

$$\begin{aligned}
\mu_0(A \cap \mathbb{M}_0) &= \frac{1}{\text{intensity}\,\mu}\mathbb{E}_\mu\left[\#\{g \in U \mid g^{-1}\omega \in A \cap \mathbb{M}_0\}\right] && \text{By definition}\\
&= \frac{1}{\text{intensity}\,\mu}\mathbb{E}_\mu\left[|\omega \cap U|\mathbb{1}_{\omega\in A}\right] && \text{By shift invariance of } A\\
&= \mu(A) && \text{By ergodicity.}
\end{aligned}$$

For the general case, we appeal to the ergodic decomposition theorem (see [Gre00] for a proof):

Theorem 13. Let G be an lcsc group, and $G \curvearrowright (X, \mu)$ a pmp action on a standard Borel space. Then there exists a standard Borel space Y equipped with a probability measure ν and a family $\{p_y \mid y \in Y\}$ of probability measures p_y on X with the following properties:

1. For every Borel $A \subset X$, the map $y \mapsto p_y(A)$ is Borel, and
$$\mu(A) = \int_Y p_y(A) d\nu(y).$$
2. For every $y \in Y$, p_y is an invariant and ergodic measure for the action $G \curvearrowright (X, p_y)$,
3. If $y, y' \in Y$ are distinct, then p_y and p'_y are mutually singular.

There is an almost identically stated version of the above theorem for pmp cbers as well.

If (Y, ν) and $\{p_y \mid y \in Y\}$ is the ergodic decomposition for $G \curvearrowright (\mathbb{M}, \mu)$, then the Palm measures $(p_y)_0$ of the p_y form an ergodic decomposition for $(\mathbb{M}_0, \mathcal{R}, \mu_0)$. That is, for all $A \subseteq \mathbb{M}_0$ we have
$$\mu_0(A) = \int_Y (p_y)_0(A) d\nu(y).$$
Applying the previous ergodic case to this yields the general formula. □

Remark 17. It is immediate that the ergodic decomposition for $G \curvearrowright (\mathbb{M}, \mu)$ determines the ergodic decomposition for $(\mathbb{M}_0, \mathcal{R}, \mu_0)$.

In the other direction, let $\{p'_y \mid y \in Y'\}$ denote the ergodic decomposition of $(\mathbb{M}_0, \mathcal{R}, \mu_0)$, so that
$$\mu_0(A) = \int_{Y'} p'_y(A) d\nu'(y).$$
It turns out that all of the ergodic components p'_y are not just probability measures on $\mathbb{M}_0$, but are themselves the Palm measures of point processes. This can be proven by using a characterisation of Mecke, see Theorem 13.2.VIII of [DVJ07] (one applies the formula listed as item (iii) to support(p'_y)).

One can then use the Voronoi inversion technique as referenced in Remark 14 to construct the ergodic decomposition of μ out of the ergodic decomposition of μ_0 (with an additional $\texttt{Unif}[0, 1]$ random variable).

1.3.3 The correspondences in the measured category

We have seen that Borel thinnings $\theta : \mathbb{M} \to \mathbb{M}$ correspond exactly to Borel subsets $A \subseteq \mathbb{M}_0$. We now wish to extend this to the measured case:

Proposition 6. There is a one to one correspondence between μ-thinnings $\theta : (\mathbb{M}, \mu) \to \mathbb{M}$ and subsets $A \subseteq (\mathbb{M}_0, \mu_0)$.

Proof. We have already described how to induce a thinning from a subset $A \subseteq (\mathbb{M}_0, \mu)$. It is the other direction where the difficulty lies.

Suppose $\theta : (\mathbb{M}, \mu) \to \mathbb{M}$ is a thinning. Then we wish to *restrict* θ to $\mathbb{M}_0$, but *this makes no formal sense a priori*. Note that $\mathbb{M}_0 \subseteq \mathbb{M}$ is a set of μ measure zero!

Nevertheless, equivariance allows us to overcome this difficulty. Observe that, by assumption,

$$\{\omega \in \mathbb{M} \mid \theta(\omega) \subseteq \omega\} \text{ has } \mu \text{ measure one.}$$

This is a shift invariant set, so by Proposition 5 we have

$$\{\omega \in \mathbb{M}_0 \mid \theta(\omega) \subseteq \omega\} \text{ has } \mu_0 \text{ measure one.}$$

Thus we have associated a subset of $(\mathbb{M}_0, \mu)$ to our μ-thinning. It is immediate that these associations are mutual inverses. □

The identical method proves

Proposition 7. There is a one to one correspondence between directed factor graphs $\mathscr{G} : (\mathbb{M}, \mu) \to \text{Graph}(G)$ and subsets of $(\overrightarrow{\mathbb{M}_0}, \overrightarrow{\mu_0})$.

1.3.4 Voronoi tesellations

Definition 18. Let $\omega \in \mathbb{M}$ be a configuration, and $g \in \omega$ one of its points. The associated *Voronoi cell* is

$$V_\omega(g) = \{x \in G \mid d(x, g) \leq d(x, h) \text{ for all } h \in \omega\}.$$

The associated *Voronoi tessellation* is the ensemble of closed sets $\{V_\omega(g)\}_{g \in \omega}$.

Left-invariance of the metric d implies that the Voronoi cells are equivariant in the sense that for all $\gamma \in G$, we have $V_{\gamma\omega}(\gamma g) = \gamma V_\omega(g)$.

Note that discreteness of the configuration implies that the Voronoi tessellation forms a locally finite *cover* of the ambient space by closed sets. We would like to think of these sets as forming a *partition* of the ambient space, but this isn't necessarily true even in the measured sense: the boundaries of the Voronoi cells can have positive volume. For example, let Γ be a discrete group and consider $\Gamma \times \{0\} \subset \Gamma \times \mathbb{R}$.

Lie groups and Riemannian symmetric spaces avoid this deficiency, as hyperplanes[7] have zero volume.

So depending on the examples one is interested in, there's not much loss of generality in thinking of the boundaries of the Voronoi cells as being Haar-null. Besides, one can sidestep this issue by introducing a *tie breaking* function. This is a Borel isomorphism $T : G \to \mathbb{R}$. Let us define

$$V_\omega^T(g) = \{x \in G \mid d(x, g) \leq d(x, h) \text{ for all } h \in \omega, \text{ and for all } h \in \omega \setminus \{g\}, T(x^{-1}g) < T(x^{-1}h)\}.$$

Note that these tie-broken Voronoi cells form a *measurable* partition of G. That is, we have traded the Voronoi cells being closed for them being genuinely disjoint. The equivariance property $V_{\gamma\omega}^T(\gamma g) = \gamma V_\omega^T(g)$ still holds as well.

1.3.5 The replication trick and factoring onto a Poisson

Lemma 1. Suppose Π and Υ are point processes, and Π factors onto Υ. Then $[0,1]^\Pi$ factors onto $[0,1]^\Upsilon$.

[7]sets of the form $\{x \in X \mid d(x, g) = d(x, h)\}$ for a fixed distinct pair $g, h \in X$

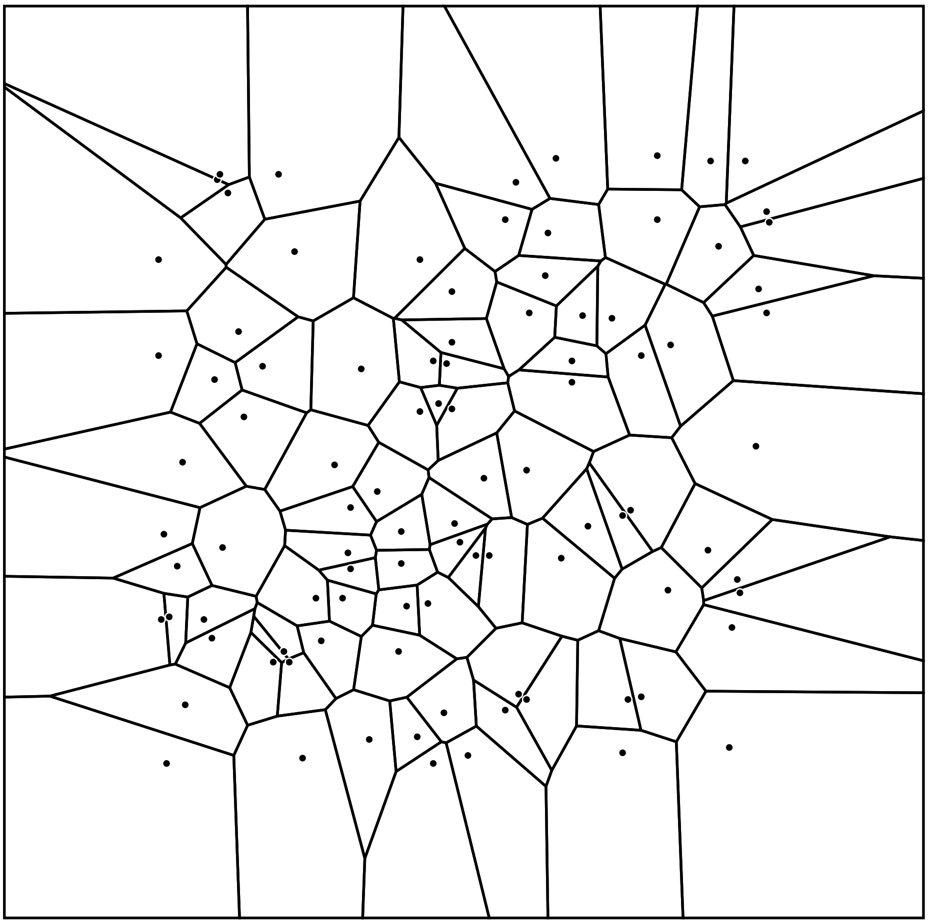

Figure 1.3: The Voronoi cells of a Poisson point process sampled on $[-10, 10]^2$ in $\mathbb{R}^2$. Note the distortion effect at the boundary: the Voronoi cells of a Poisson on $\mathbb{R}^2$ restricted to $[-10, 10]^2$ will not look like this.

The proof of this uses the following *replication trick*: that the randomness in one $\texttt{Unif}[0, 1]$ random variable ξ is equivalent to the randomness in an entire IID sequence $\xi_1, \xi_2, \dots$ of $\texttt{Unif}[0, 1]$ random variables.

More precisely, there is an isomorphism[8] (as measure spaces)

$$I : ([0, 1], \texttt{Leb}) \to ([0, 1]^{\mathbb{N}}, \texttt{Leb}^{\otimes \mathbb{N}}).$$

So if $\xi \sim \texttt{Unif}[0, 1]$, then we will write $I(\xi) = (\xi_1, \xi_2, \dots)$ for the associated IID sequence of $\texttt{Unif}[0, 1]$ random variables.

Proof of Lemma 1. Suppose $\Upsilon = \Phi(\Pi)$. If $g \in [0, 1]^{\Pi}$, then we write ξ^g for its label, and by the replication trick $\xi^g_1, \xi^g_2, \dots$ for the associated IID seqence of $\texttt{Unif}[0, 1]$ random variables.

We define a factor map Ψ of $[0, 1]^{\Pi}$ as follows:

$$\Psi([0, 1]^{\Pi}) = \bigcup_{g \in [0,1]^{\Pi}} \{(h_i, \xi^g_i) \mid V_g(\Pi) \cap \Phi(\Pi) = (h_1, h_2, \dots,)\},$$

[8]One can make the isomorphism as explicit as one wishes, but it will not aid in understanding

where we mean that $(h_1, h_2, \ldots)$ is any *enumeration* of $V_g(\Pi) \cap \Phi(\Pi)$ performed in an equivariant way.

For instance, look at the elements of $h \in V_g(\Pi) \cap \Phi(\Pi)$ which are *closest* to g. Then let h_1 be the element that minimises $T(g^{-1}h)$, where $T : G \to [0,1]$ is the tie-breaking function of Section 1.3.4. Then let h_2 be the next smallest element, and so on, until you exhaust the closest elements. Then look at the batch of next closest elements and so on. One can check that this is an equivariant construction (any construction where you do the same thing at every point will be).

Then $\Psi([0,1]^\Pi) = [0,1]^\Upsilon$, as desired. □

Proposition 8. Let Π be a point process on G. Then $[0,1]^\Pi$ factors onto the Poisson point process and the IID Poisson point process.

Proof. By the previous lemma, it suffices to prove that $[0,1]^\Pi$ factors onto the Poisson point process.

To that end, fix a map $F : [0,1] \to \mathbb{M}(G)$ such that if $\xi \sim \texttt{Unif}[0,1]$, then $F(\xi)$ is a Poisson point process on G of unit intensity.

We will use the Voronoi tessellation to simply glue independent copies of the Poisson point process in each cell, resulting in a Poisson point process.

Define a factor map $\Phi([0,1]^\Pi)$ by

$$\Phi([0,1]^\Pi) = \bigcup_{g \in [0,1]^\Pi} gF(\xi^g)\big|_{V_\Pi^T(g)}.$$

Then $\Phi([0,1]^\Pi)$ is the Poisson point process. □

Remark 18. In the study of percolation theory on a discrete group Γ a fundamental fact is that the Bernoulli percolations of different densities can all be jointly defined in a *monotone* way.

On a nondiscrete group G, it is clear how to jointly define the Poissons of lower intensity: one uses independent p-thinning, just as in the discrete case. A variation on the construction in the previous proposition shows that it is possible to jointly define the higher intensity Poissons too. This can be done in a monotone way by retaining the original set.

To sum up: if $[0,1]^\Pi$ is the unit intensity Poisson point process on G, then there exists factor maps $\Phi_t : [0,1]^\mathbb{M} \to \mathbb{M}$ indexed by $t > 0$ such that

- $\Phi_t([0,1]^\Pi)$ has the distribution of a Poisson point process of intensity t, and
- if $s < t$, then $\Phi_s([0,1]^\Pi) \subset \Phi_t([0,1]^\Pi)$.

Example 18 (The Palm version of a general thickening)**.** The Voronoi tessellation can be used to express the Palm version of an *arbitrary* thickening. Recall the setup of Example 17. Given a thickening $\Theta : \mathbb{M} \to \mathbb{M}$, define for $g \in \omega$

$$F_\omega(g) = V_\omega(g) \cap \Theta(g).$$

Then

$$\Theta(\omega) = \bigcup_{g \in \omega} gF_\omega(g),$$

and thus every thickening is of this form. One also has to ensure that these sets $F_\omega(g)$ are finite, but that follows by mass transport.

1.4 The cost of a point process

If $\mathscr{G}$ is a directed factor graph, then its *inverse* is $\mathscr{G}^{-1} = \{(g^{-1}\omega, g^{-1}) \in \overrightarrow{\mathbb{M}_0} \mid (\omega, g) \in \mathscr{G}\}$ – that is, the same graph with all the arrows reversed.

The *identity graph* is $\mathscr{I} = \mathbb{M}_0 \times \{0\}$ – that is, the graph which consists of a single loop at each vertex.

The groupoid *generated* by a factor graph $\mathscr{G}$ is

$$\langle \mathscr{G} \rangle = \bigcup_{n=0}^{\infty} (\mathscr{G} \cup \mathscr{G}^{-1} \cup \mathscr{I})^n.$$

Note that the groupoid generated by $\mathscr{G}$ is $\overrightarrow{\mathbb{M}_0}$ if and only if $\mathscr{G}$ is *connected* as an undirected graph. We refer to factor graphs with this property as *graphings*.

Definition 19. Let Π be a point process on G (possibly marked) with finite but non-zero intensity. Its *groupoid cost* is defined by

$$\text{cost}(\Pi) - 1 = \text{intensity}\, \mu \cdot \inf_{\mathscr{G}} \left\{ \mathbb{E}\left[\overrightarrow{\deg_0} \mathscr{G}(\Pi_0) \right] - 1 \right\},$$

where the infimum is taken over all graphings of Π. Equivalently by Remark 16,

$$\text{cost}(\Pi) - 1 = \inf_{\mathscr{G}} \left\{ \mathbb{E}\left[\sum_{x \in U \cap \Pi} \deg_x \mathscr{G}(\Pi) \right] \right\} - \text{intensity}(\Pi).$$

Example 19. If Π is the lattice shift corresponding to $\Gamma < G$, then

$$\text{cost}(\Pi) = 1 + \frac{d(\Gamma) - 1}{\text{covol}(G/\Gamma)},$$

where $d(\Gamma)$ denotes the *rank* of Γ, that is, its minimum number of generators.

A group is said to have *fixed price* if all of its *essentially free* point processes have the same cost. At the time of writing there are no groups known that do not have this property.

Example 20. We will later see that the any free point process Π on $\mathbb{R}^n$ is "hyperfinite", thereby giving an indirect proof that $\text{cost}(\Pi) = 1$. We will give a separate indirect proof of this fact by using the product structure: that $\mathbb{R}^n = \mathbb{R}^{n-1} \times \mathbb{R}$ for $n > 1$. Of course, it is immediate that the cost of any free process on $\mathbb{R}$ is one, so it is the $n > 1$ case we are interested in.

It would be interesting to see a more explicit proof that for every free point process on Π on $\mathbb{R}^n$ we have $\text{cost}(\Pi) < 1 + \varepsilon$ for all $\varepsilon > 0$.

Lemma 2. Let Π be a point process of finite intensity, and Φ a factor map of Π such that $\Phi(\Pi)$ has finite intensity. Then

$$\text{cost}(\Pi) \leq \text{cost}(\Phi(\Pi)).$$

Thus cost is *monotone* for factors.

Corollary 1. If μ and ν are finite intensity point processes that factor onto eachother, then $\mathrm{cost}(\mu) = \mathrm{cost}(\nu)$. In particular, the cost of μ only depends on its isomorphism class as an action.

Proof of lemma. Recall from Remark 10 that Φ decomposes as the composition of a thinning π and a thickening Θ^Φ. We prove

$$\mathrm{cost}(\Pi) \leq \mathrm{cost}(\Theta^\Phi(\Pi)) \leq \mathrm{cost}(\pi(\Theta^\Phi(\Pi)) = \mathrm{cost}(\Phi(\Pi)),$$

where the last equality holds as $\Phi = \pi \circ \Theta^\Phi$.

We prove the second inequality first, as it is simpler. For this we use the non-Palm definition of cost.

To that end, let $\mathscr{G}$ be a graphing of $\Phi(\Pi)$ that ε-computes the cost, that is, with

$$\mathbb{E}\left[\sum_{x \in U \cap \Phi(\Pi)} \overrightarrow{\deg}_x \mathscr{G}(\Phi(\Pi))\right] - \mathrm{intensity}(\Phi) \leq \mathrm{cost}(\Phi(\Pi)) - 1 + \varepsilon.$$

We will use it to define a graphing $\mathscr{H}$ of the thickened process $\Theta^\Phi(\Pi)$. Recall that this process has three types of points: red, purple, and blue.

Let $\mathscr{N}$ be the factor graph of $\Theta^\Phi(\Pi)$ that connects each red point x to its nearest blue neighbour. If this is not well-defined, then we use the tie-breaking function $T : G \to \mathbb{R}$ of Section 1.3.4 to make it so in an equivariant way.

That is, if $y_1, y_2, \ldots, y_n$ are the (finitely many!) blue points of $\Theta^\Phi(\Pi)$ that are closest to x, then let y be the element that minimises $T(x^{-1}y_i)$ and add in a directed edge $x \to y$ to $\mathscr{N}$.

We can view $\mathscr{G}$ as defining a factor graph on $\Theta^\Phi(\Pi)$, which lives on the blue and purple points.

Now let $\mathscr{H}(\Theta^\Phi(\Pi) = \mathscr{G}(\Phi(\Pi)) \sqcup \mathscr{N}(\Theta^\Phi(\Pi)$. This is connected as an undirected graph, so by definition of cost:

$$\begin{aligned}
&\mathrm{cost}(\Theta^\Phi(\Pi)) - 1 \leq \mathbb{E}\left[\sum_{x \in \Theta^\Phi(\Pi) \cap U} \overrightarrow{\deg}_x \mathscr{H}(\Theta^\Phi(\Pi))\right] - \mathrm{intensity}(\Theta^\Phi(\Pi)) \\
&= \mathbb{E}\left[\sum_{x \in U \cap \Pi \setminus \Phi(\Pi)} 1 + \sum_{x \in U \cap \Phi(\Pi)} \overrightarrow{\deg}_x \mathscr{G}(\Phi(\Pi))\right] - \mathrm{intensity}(\Pi \setminus \Phi) - \mathrm{intensity}(\Phi) \\
&= \mathbb{E}\left[\sum_{x \in U \cap \Phi(\Pi)} \overrightarrow{\deg}_x \mathscr{G}(\Phi(\Pi))\right] - \mathrm{intensity}(\Phi) \\
&\leq \mathrm{cost}(\Phi(\Pi)) - 1 + \varepsilon.
\end{aligned}$$

As ε was arbitrary, this proves the second inequality.

For the other inequality, we use the explicit description of the Palm measure as in Example 18 and the Palm definition of cost.

The idea of the proof is: we have a graphing defined on a larger subset, and we must push it onto a smaller subset somehow. We simply transfer all edges of $\Theta^\Phi(\Pi)$ to Π along the Voronoi cells.

For $g \in \Pi$, let $F_\Pi(g) = V_\Pi(g) \cap \Theta^\Phi(\Pi)$.

Let us call a graphing $\mathscr{G}$ of $\Theta^\Phi(\Pi)$ *starlike* if it contains if for all $g \in \Pi$ and $x \in F_\Pi(g)$, we have $(g, x) \in \mathscr{G}$. If $\mathscr{G}$ is any graphing, then we can perturb it to find a starlike graphing of the same edge measure.

Let $\mathscr{G}$ be a starlike graphing of $\Theta^\Phi(\Pi)$ that ε-computes the cost. Let us define a graphing $\mathscr{H}$ of Π as follows: join $x, y \in \Pi$ by an edge in $\mathscr{H}(\Pi)$ if there exists $x' \in F_\Pi(x)$ and $y' \in F_\Pi(y)$ such that x' and y' are connected by an edge in $\Theta^\Phi(\Pi)$.

When we push $\mathscr{G}$ onto Π, some edges get killed. Obviously if two Voronoi cells have many edges between them, then some get killed. But note by the *starlike* assumption, every edge *within* the Voronoi cell gets killed too. In particular, we kill $|F_\Pi(g)| - 1$ edges at each $g \in \Pi$.

To make the proof more legible, let us write $I_\Pi = \text{intensity}(\Pi)$ and $I_\Theta = \text{intensity}(\Theta^\Phi(\Pi))$, so that $I_\Theta = I_\Pi \cdot \mathbb{E}[F_{\Pi_0}(0)]$.

We compute its expected outdegree as follows:

$$I_\Pi \cdot \mathbb{E}\left[\overrightarrow{\deg}_0 \mathscr{H}(\Pi_0) - 1\right] \leq I_\Pi \cdot \mathbb{E}\left[\sum_{x \in F_{\Pi_0}(0)} \overrightarrow{\deg}_x \mathscr{G}(\Theta^\Phi(\Pi_0)) - |F_{\Pi_0}(0)|\right]$$

$$= I_\Pi \cdot \mathbb{E}\left[\sum_{x \in F_{\Pi_0}(0)} \overrightarrow{\deg}_x \mathscr{G}(\Theta^\Phi(\Pi_0))\right] - I_\Pi \cdot \mathbb{E}|F_{\Pi_0}(0)|$$

$$= I_\Pi \cdot \mathbb{E}\left[\sum_{x \in F_{\Pi_0}(0)} \overrightarrow{\deg}_x \mathscr{G}(\Theta^\Phi(\Pi_0))\right] - I_\Theta.$$

We now work on this first term.

$$I_\Pi \cdot \mathbb{E}\left[\sum_{x \in F_{\Pi_0}(0)} \overrightarrow{\deg}_x \mathscr{G}(\Theta^\Phi(\Pi_0))\right] = \frac{I_\Theta}{\mathbb{E}|F_{\Pi_0}(0)|}\mathbb{E}\left[\sum_{x \in F_{\Pi_0}(0)} \overrightarrow{\deg}_x \mathscr{G}(\Theta^\Phi(\Pi_0))\right]$$

$$= \sum_{k \geq 1} \frac{I_\Theta}{\mathbb{E}|F_{\Pi_0}(0)|}\mathbb{E}\left[\sum_{x \in F_{\Pi_0}(0)} \overrightarrow{\deg}_x \mathscr{G}(\Theta^\Phi(\Pi_0)) \Big| |F_{\Pi_0}(0)| = k\right] \mathbb{P}[|F_{\Pi_0}(0)| = k]$$

$$= I_\Theta \sum_{k \geq 1} \mathbb{E}\left[\frac{1}{|F_{\Pi_0}(0)|} \sum_{x \in F_{\Pi_0}(0)} \overrightarrow{\deg}_x \mathscr{G}(\Theta^\Phi(\Pi_0)) \Big| |F_{\Pi_0}(0)| = k\right] \mathbb{P}[|F_{\Pi_0}(0)| = k]$$

$$= I_\Theta \mathbb{E}\left[\overrightarrow{\deg}_X \mathscr{G}(\Theta^\Phi(\Pi_0))\right]$$

$$= I_\Theta \mathbb{E}\left[\overrightarrow{\deg}_0 (\mathscr{G}(\Theta^\Phi(\Pi)_0))\right]$$

Thus

$$I_\Pi \cdot \mathbb{E}\left[\overrightarrow{\deg}_0 \mathscr{H}(\Pi_0) - 1\right] \leq I_\Theta \mathbb{E}\left[\overrightarrow{\deg}_0 (\mathscr{G}(\Theta^\Phi(\Pi)_0)) - 1\right],$$

proving $\text{cost}(\Pi) \leq \text{cost}(\Theta^\Phi(\Pi))$, as desired. □

Remark 19. The groupoid cost can really increase under a factor map: take the example of Remark 9 with $\mathbb{Z}^n < \mathbb{R}^n$ for $n > 1$.

1.4.1 Cost is finite for compactly generated groups

Proposition 9. Suppose G is *compactly generated* by $S \subseteq G$. Then every free point process μ on G has finite cost (implicitly we are assuming μ has finite intensity, as it must for the cost to even be defined).

We recall some definitions and facts from metric geometry, see [CdlH16] for further details in the specific context we are interested in.

Definition 20. Let (X, d) be a metric space.

- (X, d) is *coarsely connected* if there exists $c > 0$ such that for all $x, x' \in X$ there are points $x_1, x_2, \ldots, x_n \in X$ with $x = x_1$, $x_n = x'$, and $d(x_i, x_{i+1}) \leq c$ for all i.
- A subset $\omega \subseteq X$ is *uniformly discrete* if there exists $\varepsilon > 0$ such that $d(x, y) > \varepsilon$ for all distinct $x, y \in \omega$.
- A subset $\omega \subseteq X$ is *coarsely dense* if there exists $r > 0$ such that for every $x \in X$, $d(x, \omega) < r$.
- A *Delone set* is a subset $\omega \subseteq X$ which is both uniformly discrete and coarsely dense.
- An *ε-net* is a subset $\omega \subseteq X$ which is $\frac{\varepsilon}{2}$ uniformly discrete and ε coarsely dense.

Theorem 14 (See Proposition 1.D.2 of [CdlH16])**.** Let G be an lcsc group with a left-invariant proper metric d which generates its topology. Then G is compactly generated if and only if it is coarsely connected and large-scale geodesic.

Note that if X is coarsely connected, then so to is any coarsely dense subset of X.

Definition 21. Let $S \subseteq G$ be a compact generating set.

The *Cayley factor graph* associated to S is the map $\text{Cay}(\bullet, S) : \mathbb{M} \to \text{Graph}(G)$ given by

$$\text{Cay}(\omega, S) = \{(g, gs) \in \omega \times \omega \mid s \in S\}.$$

Note that this graph is not necessarily connected. However, if Π is a point process which is almost surely R-coarsely-dense for $R < \operatorname{diam} S$, then $\text{Cay}(\Pi, S)$ is connected. This condition can always be satisfied by choosing an appropriate power S^k of the generating set S.

The following can be readily deduced from existing results in the literature (even removing the compact generation assumption), but we include a separate proof for completeness and ideological reasons.

Proposition 10. Suppose Π is a free point process on a compactly generated group G. Then for every $R > 0$ there exists a *finite intensity* thickening Θ of Π such that $\Theta(\Pi)$ is almost surely R-coarsely-dense.

Moreover, if Π is δ-separated (with $\delta < 2R$), then Θ will also be δ-separated.

Proof. Fix $R > 0$. We will construct a factor map Φ of Π such that $\Phi(\Pi)$ is $\frac{R}{2}$ uniformly separated and $\Theta(\Pi) := \Pi \sqcup \Phi(\Pi)$ is R-coarsely dense. The uniform separation implies then implies that this thickening has finite intensity.

The idea of the proof is the following: observe that every uniformly separated subset of a metric space is a subset of a Delone set. You can prove this using the well-ordering principle or Zorn's lemma (as to your taste). Now consider a sample Π from the point process. We know there are *some* ways to add points to it to get something coarsely dense, the only difficulty is that we are required to make these choices equivariantly. We will select points that see the "frontier" of the process, which will then add points to cover a piece of the frontier. At every stage the frontier gets smaller, and in the limit we cover the whole space.

For configurations $\omega \in \mathbb{M}$, let ω^t denote[9]

$$\omega^t = \bigcup_{g \in \omega} B(g,t)$$

the union of all *closed* balls about the points of ω.

We call a point $g \in \Pi$ *on the frontier* if $B(g,2R) \not\subseteq \Pi^R$, and let $F(\Pi)$ denote the subset of frontier points of Π. This is a metrically defined condition, and hence equivariant. We will define a rule $\Phi_1(\Pi)$ that specifies a collection of points such that their R-balls cover all the $2R$-balls of the frontier points of Π. We will then iterate this construction (so that $\Phi_2(\Pi)$'s R-balls cover the $2R$-balls of $\Phi_1(\Pi) \cup \Pi$'s frontier points, and so on). In this way we will find enough points to cover the whole space[10].

One can decompose the frontier points of Π as

$$F(\Pi) = \bigsqcup_n F_n(\Pi),$$

where each $F_n(\Pi)$ is $10R$ uniformly separated. This can be done by using the existence of a *Borel kernel* of the factor graph $\mathscr{D}_{10R}$ defined on the frontier points, see Section 4 of [KST99] for further information.

We now fix an auxiliary (deterministic) R-net $\mathcal{N} \subset G$. If $W \subseteq G$ is a Borel region and $g \in G$, then let

$$N(g,W) = \{x \in g^{-1}\mathcal{N} \mid B(x,R) \cap W \neq \varnothing\}.$$

Note that $N(g,W)^R \supseteq W$, as $\mathcal{N}$ is coarsely dense. Define

$$\Phi_1(\Pi) = \bigcup_{g \in F_1(\Pi)} N(g, B(g,2R) \setminus \Pi^R),$$

and inductively

$$\Phi_{n+1}(\Pi) = \bigcup_{g \in F_n(\Pi)} N(g, B(g,2R) \setminus \left(\Pi \cup \bigcup_{i \leq n} \Phi_i(\Pi)\right)^R)$$

Then

$$\Pi^{2R} \subseteq \Pi^R \cup \bigcup_{n \geq 1} \Phi_n(\Pi)^R.$$

Iterating this procedure allows you to equivariantly construct an increasing sequence of uniformly separated factors that cover Π^{nR} for all $n \in \mathbb{N}$. The union of these is the desired factor.

□

[9] We are defining the most simple random closed set associated to Π, a particular example of something called *the Boolean model*. This is supposed to reassure the finicky that the measure theoretic details can be worked out.

[10] Read Settlers [Sak14].

Remark 20. We will later describe the connection between point processes and "cross-sections" of actions. The previous proposition can be deduced from the fact that every free action admits a "cocompact cross-section". A similar statement to the proposition directly phrased in terms of cross-sections can be found in Section 2 of [Slu17], where it is shown that any cross-section can be extended to a cocompact cross-section. That proof works without the compact generation assumption.

1.5 Amenability

In this section, we will characterise amenability of a group in terms of the free point processes on it. Whilst not especially novel, this will clarify certain results in the literature. For our purposes the most convenient definition of amenability will be the existence of a left-invariant mean $m \in (L^\infty(G))^*$.

Holroyd and Peres introduced the following concept in [HP03]:

Definition 22. Let Π be a point process with law μ. A sequence of factor graphs $\sim_n^\bullet \colon (\mathbb{M}, \mu) \to \mathrm{Graph}(G)$ is a *one-ended clumping* if it satisfies the following for μ almost every $\omega \in \mathbb{M}$:

- (Ascending) $\sim_1^\omega \subseteq \sim_2^\omega \subseteq \cdots$
- (Partitions) the connected components of each $\sim_n^\omega$ consist of *finite* complete graphs, and
- (One-endedness) for all x, y in ω there exists $N = N(x, y, \omega)$ such that x is connected to y in $\sim_N^\omega$.

We view $\sim_n^\omega$ as an *equivalence relation* on ω consisting of finite classes. If $x, y \in \omega$ then we will write $x \sim_n^\omega y$ to denote that x and y are connected in $\sim_n^\omega$.

Recall that if Π is a point process, then the ensemble of Voronoi cells $\{V_\Pi(g)\}_{g \in \Pi}$ forms a random measurable partition of G. If $\sim_n^\bullet$ is a clumping of Π, then it gives us a way to *coarsen* the Voronoi partitioning as follows: for each n, define

$$\mathcal{P}_n = \left\{ \bigcup_{h \sim_n^\Pi g} V_\Pi(h) \right\}_{g \in \Pi}.$$

Note that $\mathcal{P}_n$ is a refinement of $\mathcal{P}_{n+1}$. See Figure 1.5.

Holroyd and Peres were interested in (among other things) constructing particular kinds of connected factor graphs on the Poisson point process. Namely, they were interested in constructing one-ended factor trees and directed $\mathbb{Z}$s. They proved:

Theorem 15 (Holroyd-Peres [HP03])**.** Let Π denote a free and ergodic point process in $\mathbb{R}^n$. Then the following are equivalent:

- Π admits a locally finite factor graph which is a connected and one-ended tree,
- Π admits a factor graph which is isomorphic to the directed line $\mathbb{Z}$, and
- Π admits a one-ended clumping.

Moreover, the Poisson point process admits a one-ended clumping.

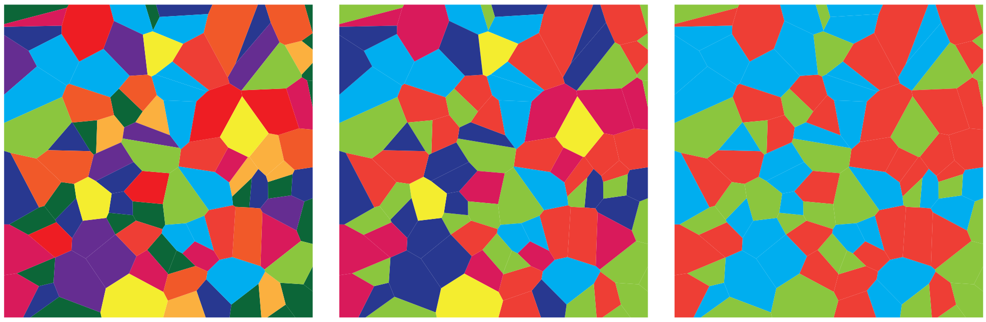

$\mathcal{P}_1$ $\mathcal{P}_2$ $\mathcal{P}_3$

[h]

Figure 1.4: This is how you should visualise the partitions associated to a one-ended clumping: like a sequence of worse and worse mosaics.

This was later extended by Ádám Timár, who also answered a question of Steve Evans about possible factor graph structures on point processes:

Theorem 16 (Timár [Tim04])**.** Let Π denote a free and ergodic point process in $\mathbb{R}^n$. Then Π admits a one-ended clumping. Moreover, Π admits a connected factor graph isomorphic to $\mathbb{Z}^d$, for any $d \in \mathbb{N}$.

It is clear from these works that the amenability of the underlying space $\mathbb{R}^n$ is important, but the connection was not fully elucidated. We will prove

Theorem 17. If G is amenable, then all of its free point processes admit one-ended clumpings. Conversely, if G has a free point process that admits a one-ended clumping, then G is amenable.

The same is true for marked point processes.

As a statement in its own right, this isn't that impressive (and especially with the framing we've given it, where it seems inevitable). The *real* statement is that all of the factor graph related questions are governed by the associated Palm equivalence relation.

Recall the following:

Definition 23. A pmp cber $(X, \mathcal{R}, \mu)$ is *μ-hyperfinite* if there exists an increasing sequence $\mathcal{R}_1 \subseteq \mathcal{R}_2 \subseteq \cdots$ of subequivalence relations of $\mathcal{R}$ such that for μ almost every $x \in X$,

- for all $n \in N$, $[x]_{\mathcal{R}_n}$ is finite, and
- $[x]_{\mathcal{R}} = \bigcup_n [x]_{\mathcal{R}_n}$.

Denote by $\widetilde{\mu}$ the lifted measure of μ to $X \times X$.

A pmp cber $(X, \mathcal{R}, \mu)$ is *μ-amenable* if there exists for each $x \in X$ a normalised positive functional $p_x \in (\ell^\infty([x]_{\mathcal{R}}))^*$ (a *local mean*) such that $p_x = p_y$ for $\widetilde{\mu}$ almost every $(x, y) \in X \times X$, and such that the function $x \mapsto p_x\left(\varphi|_{[x]_{\mathcal{R}}}\right)$ is measurable for all $\varphi \in L^\infty(X, \mu)$.

In the measured category, these concepts are equivalent (see Chapter II section 10 of [KM04]):

Theorem 18 (Connes-Feldman-Weiss)**.** A pmp cber $(X, \mathcal{R}, \mu)$ is μ-hyperfinite if and only if it is μ-amenable.

Under the correspondences we've described, a free point process Π admits a one-ended clumping *if and only if* its Palm equivalence relation $(\mathbb{M}_0, \mathcal{R}, \mu_0)$ is μ_0-hyperfinite.

Proof of Theorem 17. The proof will be the same for marked and unmarked processes, so we work with unmarked ones for notational convenience.

We first describe how a one-ended clumping can be used to construct an invariant mean on G using a fairly standard technique, see Theorem 5.1 of [BLPS99].

Let μ be a free point process, and fix a clumping $(\sim_n^\bullet)$ of it. If $f : G \to \mathbb{R}$ is an essentially bounded function, define

$$m_n(f) = \frac{1}{\text{intensity}\, \mu} \mathbb{E}_\mu \left[\sum_{y \sim_n^\omega X(\omega)} \frac{f(y)}{\#\{y \sim_n^\omega X(\omega)\}} \right],$$

that is, we average the values of f over the points in $X(\omega)$'s nth equivalence class.

By invariance of the point process, one can see that

$$m_n(g \cdot f) = \frac{1}{\text{intensity}\, \mu} \mathbb{E}_\mu \left[\sum_{y \sim_n^\omega X^g(\omega)} \frac{f(y)}{\#\{y \sim_n^\omega X^g(\omega)\}} \right].$$

One-endedness of the clumping implies that for n sufficiently large, $\{y \sim_n^\omega X(\omega)\} = \{y \sim_n^\omega X^g(\omega)\}$. So any ultralimit of the m_n defines a left-invariant mean on G.

Fix a left-invariant mean $m \in (L^\infty(G))^*$. Given a bounded and positive function $f : [\omega]_\mathcal{R} \to \mathbb{R}$, we extend it to a function $F : G \to \mathbb{R}$ by making F constant on the Voronoi cells, and averaging the values for those $g \in G$ that belong to multiple Voronoi cells[11]. Define

$$F(g) = \sum_{\{x \in \omega \mid g \in V_\omega(x)\}} \frac{f(x^{-1}\omega)}{\#\{x \in \omega \mid g \in V_\omega(x)\}}.$$

Now for each $\omega \in \mathbb{M}_0$ we define a mean on $[\omega]_\mathcal{R}$ by $p_\omega(f) = m(F)$. Then p_ω only depends on the equivalence class of ω by left-invariance of m, and satisfies the measurability requirement. □

Remark 21. A version of this theorem was independently proved by Paquette in [Paq18]. He looks specifically at invariant point processes on Riemannian symmetric spaces and (among other things) proves that the Delauney triangulation of any point process on such a space is a unimodular random network which is *anchored amenable* if and only if the ambient space is amenable.

1.6 Factor constructions

1.6.1 Independent sets

The following is a tool that we use repeatedly. There are multiple ways to prove it, the following proof appears in [Tim04] and is attributed to Yuval Peres.

[11]Note that each point only belongs to finitely many Voronoi cells, by local finiteness of the configuration

Lemma 3. Let μ be a *free* point process on G, and $\mathscr{G}$ a locally finite factor graph of μ. Then one can equivariantly construct a non-trivial independent subset of $\mathscr{G}$.

To spell this out, this means there exists a map $I : (\mathbb{M}, \mu) \to \mathbb{M}$ with the properties that

- $I(\omega) \subset \omega$ almost surely, and
- if $g, h \in I(\omega)$, then g and h are not connected in $\mathscr{G}(\omega)$.

Proof. Consider the factor labelling $\odot : \mathbb{M} \to \mathbb{M}_0{}^{\mathbb{M}}$ given by

$$\odot(\omega) = \{(g, g^{-1}\omega) \in G \times \mathbb{M}_0(G) \mid g \in \omega\}.$$

Under $\odot$, each point g of a configuration ω looks at how the configuration looks like from its perspective, and records it as a label. That is, it views itself as the centre of the universe (this is what the symbol $\odot$ is meant to represent, we will call the map *egotistical* or *self-centred*).

The key observation is that μ is an (essentially) free action if and only if $\odot(\omega)$ has distinct labels almost surely. For if $g, h \in \omega$ receive the same label under the egotistical map, then $g^{-1}\omega = h^{-1}\omega$, ie. $gh^{-1} \in \text{stab}_G(\omega)$. Conversely, if $g \in \text{stab}_G(\omega)$ is nontrivial, then for all $x \in \omega$ the label $x^{-1}\omega$ of x is the same as that of gx, as $(gx)^{-1}\omega = x^{-1}\omega$.

Fix a countable dense subset $Q \subset \mathbb{M}_0$. Let us define a thinning $I_q : \mathbb{M} \to \mathbb{M}$ for each $q \in Q$ by

$$I_q(\omega) = \{g \in \omega \mid d(g^{-1}\omega, q) < d(h^{-1}\omega, q) \text{ for all } h \in \omega \text{ adjacent to } g \text{ in } \mathscr{G}(\omega)\}.$$

Note that each $I_q(\omega)$ is an independent subset of $\mathscr{G}(\omega)$, but it is possibly empty. However, the *union* over all q of the I_q is ω by freeness, so at least one such I_q must define a non-empty independent subset, as desired.

□

In particular, by applying the lemma to the factor graph $\mathscr{D}_R$ of Example 10, one has:

Corollary 2. Let Π be a free point process. Then for all $R > 0$ one can deterministically and equivariantly select a subset $\Pi_R \subset \Pi$ that is R uniformly separated, in the sense that if x and y are distinct points of Π_R, then $d(x, y) > R$.

The same egotistical map will be used in the proof of the following hack, which we will need later. We give it a disparaging name as it is not to be respected.

Proposition 11 (Label trickery)**.** Let Π be any free point process, and $\theta(\Pi)$ any nonempty thinning. Then there exists a marked point process Υ such that the underlying point set of Υ is $\theta(\Pi)$, and Υ is isomorphic to Π *as a pmp action*. In particular, Υ is a free action.

The same can be achieved with Υ having marks from the *compact* space $[0, 1]$.

Proof. Let $\Upsilon = \theta(\odot(\Pi))$, that is,

$$\Upsilon = \{(g, g^{-1}\Pi) \in \mathbb{M} \times \mathbb{M}_0 \mid g \in \theta(\Pi)\}.$$

Observe that this is an *injective* map, as one can recover Π uniquely from the knowledge of any point Υ and its label, and so Υ is an isomorphic process to Π.

For the second statement, simply fix a Borel isomorphism[12] $I : \mathbb{M}_0 \to [0, 1]$, and define

$$\Upsilon = \{(g, I(g^{-1}\Pi) \in \mathbb{M} \times [0, 1] \mid g \in \theta(\Pi)\}.$$

□

[12]It exists as $\mathbb{M}_0$ is a Polish space, and thus standard Borel, and all standard Borel spaces of the same cardinality are isomorphic

1.6.2 Amenable Cayley graphs

Proposition 12. Suppose G is a unimodular amenable group, and Γ is a countably infinite amenable group, finitely generated by $S \subseteq \Gamma$. Then every free and ergodic point process μ on G admits a factor graph isomorphic to $\text{Cay}(\Gamma, S)$.

Proof. Since $(\mathbb{M}_0, \mathcal{R}, \mu_0)$ is μ_0-hyperfinite, there exists an *orbit equivalence* $\varphi : (\mathbb{M}_0, \mu_0) \to ([0,1]^\Gamma, \texttt{Leb}^{\otimes \Gamma})$, that is, a measure space isomorphism satisfying $\varphi([\omega]_\mathcal{R}) = \Gamma\varphi(\omega)$ for μ_0 almost every ω. We simply use this isomorphism to transfer the graph, using the fact that ω is bijectively equivalent with its rerooting equivalence class $[\omega]_\mathcal{R}$: define

$$\mathscr{G}(\omega) = \{(g,h) \in \omega \times \omega \mid \exists s \in S \text{ such that } \varphi(h^{-1}\omega) = \varphi(h^{-1}\omega)s\}.$$

Then $\mathscr{G}$ is the desired factor graph. □

1.6.3 Poissons on property (T) groups vs. Cayley graphs

Proposition 13. Suppose G is noncompact and has property (T), and that G has no compact normal subgroups. Then the Poisson point process on G admits no factor graph of the form $\text{Cay}(\Gamma, S)$ for *any* discrete group Γ. Furthermore, even the IID Poisson point process on G admits no such factor graph.

Equivalently: the Palm equivalence relation of the Poisson point process on such a group cannot be *freely* generated by the action of any discrete group.

We will prove this as a straightforward application of Popa's cocycle superrigidity theorem with an additional piece of point process technology called "the extra head scheme". Two great introductory sources for understanding the superrigidity theorem are Alex Furman's survey [Fur09] and the book of Kerr and Li [KL16].

Definition 24 (Malleability)**.** Let $G \curvearrowright (X, \mu)$ be a pmp action. Recall that the *weak topology* on $\text{Aut}(X, \mu)$ is the weakest topology that makes all functions $T \mapsto \mu(TA)$ continuous, where $T \in \text{Aut}(X, \mu)$ and $A \subseteq X$ is Borel.

The *flip* element of $\text{Aut}(X \times X, \mu \otimes \mu)$ is $\texttt{FLIP}(x, y) = (y, x)$.

Note that G acts on $\text{Aut}(X \times X, \mu \otimes \mu)$ diagonally via $(g \cdot T)(x, y) := T(gx, gy)$, and $\texttt{FLIP}$ commutes with this action.

The action $G \curvearrowright (X, \mu)$ is *malleable* if there exists a continuous path $\gamma : [0,1] \to \text{Aut}(X \times X, \mu \otimes \mu)$ from id to $\texttt{FLIP}$ such that γ_t commutes with the diagonal action for every $t \in [0,1]$.

The following fact seems to have gone unobserved:

Proposition 14. The IID Poisson point process is malleable.

Proof. Observe that a sample from $[0,1]^\mathscr{P} \otimes [0,1]^\mathscr{P}$ (that is, sampling from two independent unit intensity IID Poissons and keeping track of which is which) is the same as sampling from an IID Poisson Π of double the intensity with labels from $[0,1] \times \{\pm 1\}$.

Define for $0 \leq t \leq 1$ the map $t : [0,1] \times \{\pm\} \to [0,1] \times \{\pm\}$ by

$$\varphi_t(x, i) = \begin{cases} (x, -i) & x \leq t \\ (x, i) & \text{else.} \end{cases}$$

Now define

$$\gamma_t(\Pi) = \{(g, \varphi_t(x, i)) \in G \times [0,1] \times \{\pm\} \mid (g, x, i) \in \Pi\}.$$

Then (in these coordinates, γ_t continuously deforms id to $\texttt{FLIP}$. □

Definition 25. Let Γ be a discrete group and $\mathcal{G}$ a groupoid. A Γ-*valued cocycle* of the groupoid is a measurable function $c : \mathcal{G} \to \Gamma$ satisfying *the cocycle identity*

$$c(\omega, g) \cdot c(g^{-1}\omega, h) = c(\omega, gh) \text{ for all } \omega \in \mathbb{M}_0 \text{ and } g, gh \in \omega.$$

Two cocycles $c, c' : \overrightarrow{\mathbb{M}_0} \to \Gamma$ are *cohomologous* if there exists a measurable function $f : \mathbb{M}_0 \to \Gamma$ such that for all $(\omega, g) \in \overrightarrow{\mathbb{M}_0}$

$$c'(\omega, g) = f(g\omega)c(\omega, g)f(\omega).$$

Remark 22. Recall that in the categorical framework, a groupoid is a category where every arrow is invertible, and a group is the same thing but with only one object. In this language, a cocycle is a functor from a groupoid to a group, and two such cocycles are cohomologous exactly when there's a natural transformation between the two functors.

Example 21. If $G \curvearrowright (X, \mu)$ is a pmp action, then the associated *action groupoid* has unit space (X, μ) and arrows of the form (x, g) for $x \in X$ and $g \in G$. The source of such an arrow is x, and its target is gx. The composition rule for arrows is

$$(x, g) \cdot (y, h) := (x, gh) \text{ if } y = gx.$$

Note that if $\rho : G \to \Gamma$ is a homomorphism, then it induces a cocycle $c_\rho(\omega, g) = \rho(g)$. We will abuse notation and denote this cocycle simply by ρ.

In an identical way we see that ρ can be viewed as a cocycle of $\overrightarrow{\mathbb{M}_0}$.

We will use the following very basic form of Popa's cocycle superrigidity theorem:

Theorem 19. Let $G \curvearrowright (X, \mu)$ be an ergodic and weakly mixing pmp action of an lcsc group G with property (T). Then any cocycle $c : G \times X \to \Gamma$ of the action groupoid is cohomologous to a homomorphism $\rho : G \to \Gamma$.

We will describe a method of *inducing* cocycles of the rerooting groupoid to the action groupoid. This requires the following piece of point process technology:

Definition 26. Let Π be a point process on G. A (nonrandomized) *extra head scheme* is a measurable function $\mathcal{E}_\bullet : \mathbb{M} \to G$ with the property that the random variable $\mathcal{E}_\Pi^{-1}\Pi$ is distributed according to the Palm measure of Π.

An extra head scheme is thus a rule that allows one to factor a point process onto its own Palm measure in a particular way: one plucks a root (or "extra head") out of the process itself. Our interest in extra heads is their connection to orbit equivalence: under the map $\omega \mapsto \mathcal{E}_\omega^{-1}\omega$, shift equivalent configurations are mapped to rootshift equivalent configurations.

Holroyd and Peres proved the following theorem in [HP05]:

Theorem 20 (Holroyd-Peres)**.** Let Π be an ergodic point process on *nondiscrete* G. Then a nonrandomized extra head scheme always exists.

Remark 23. There is a version of the above statement for discrete groups, but to ensure that the extra head scheme is *nonrandomized* one has to further assume that the intensity is the reciprocal of an integer. This is one example of a distinction between discrete and nondiscrete groups.

Proof of proposition. The equivalence of the two statements is immediate by the correspondences we've developed, so we prove the latter.

We will denote the IID Poisson on G by μ and its Palm measure by μ_0.

The Palm equivalence relation of the IID Poisson directly factors onto that of the Poisson, so we will prove that stronger statement.

Suppose Γ acts on $([0,1]^{\mathbb{M}_0}, \mu_0)$ in a pmp and free way, generating the rootshift equivalence relation. Fix an extra head scheme $\mathcal{E}_\bullet$ for the IID Poisson. Then we can construct a cocycle $c : G \times ([0,1]^{\mathbb{M}}, \mu) \to \Gamma$ in the following way:

$$c(g, \omega) = \gamma, \text{ where } \gamma \text{ is the unique element such that } \gamma \cdot \mathcal{E}_\omega^{-1}\omega = \mathcal{E}_{g\omega}^{-1} gw.$$

(Such γ exist by the generation assumption, and are unique by freeness).

A simple computation shows that $c(g, \omega)$ is a cocycle. Thus by Popa's cocycle superrigidity we can find a homomorphism $\rho : G \to \Gamma$ and a measurable function $f : ([0,1]^{\mathbb{M}}, \mu) \to \Gamma$ such that

$$c(g, \omega) = f(g\omega)\rho(g)f(\omega)^{-1}.$$

Note that ρ is measurable, and thus by automatic continuity [citation] is in fact *continuous*. This implies $\ker \rho$ is a clopen subgroup of positive measure, and hence infinite measure by assumption on G. Equivalently, $\ker \rho$ is noncompact.

Note that for $g \in \ker \rho$, we have $c(g, \omega) = f(g\omega)f(\omega)^{-1}$. By definition of the cocycle then

$$f(g\omega)^{-1}\mathcal{E}_{g\omega}^{-1} g\omega = f(\omega)^{-1}\mathcal{E}_\omega^{-1}\omega.$$

That is, the function $\omega \mapsto f(\omega)^{-1}\mathcal{E}_\omega^{-1}\omega$ is N-invariant. The IID Poisson point process is a mixing action for G, and hence also for N by noncompactness. Therefore this N-invariant function must be *constant* by ergodicity. Note that $f(\omega)^{-1}\mathcal{E}_\omega^{-1}\omega \in [\omega]_{\mathcal{R}}$ for every ω. Thus if this function is a constant $\Omega \in \mathbb{M}_0$, we would have

$$\mathbb{P}[\Pi_0 \notin [\Omega]_{\mathcal{R}}] = 0, \text{ however } \mathbb{P}[\Pi_0 \in [\Omega]_{\mathcal{R}}] \leq \sum_{g \in \Omega} \mathbb{P}[\Pi_0 = g^{-1}\Omega] = 0,$$

a contradiction. □

Remark 24. The extra head scheme was essential in defining the cocycle c. One could try replacing it by using Voronoi cells to form a kind of discrete approximation to G: that is, let $X_g(\omega)$ be the unique element of ω such that $g \in V_\omega(X_g(\omega))$. Then define $c(g, \omega) = \gamma$, where $\gamma \cdot X_0^{-1}\omega = X_g^{-1}g\omega$. This satisfies the cocycle identity, but the issue is that $X_0^{-1}\omega$ is *not* distributed according to μ_0. This is related to *the waiting time paradox*.

1.6.4 Gaboriau-Lyons for point processes

Let $\mathbb{F}_2 = \langle a, b \rangle$ denote the free group on two generators. This is the most basic example of a nonamenable group, and consequently any discrete group containing $\mathbb{F}_2$ as a subgroup is nonamenable. There are examples of discrete groups that are nonamenable but do not contain $\mathbb{F}_2$ as a subgroup[13]. Nevertheless, every nonamenable group *measurably* contains $\mathbb{F}_2$ in the following sense:

[13] A group which is not free is *unfree*, a group with every subgroup unfree is *hereditarily unfree*, or *hunfree* for short.

Theorem 21 (Measurable von Day)**.** Let Γ be a countable nonamenable group. Then for any non-trivial[14] probability space (K, κ), there exists an ergodic and pmp action $\mathbb{F}_2 \curvearrowright (K, \kappa)^\Gamma$ such that

$$\mathbb{F}_2\omega \subseteq \Gamma\omega \text{ for almost every } \omega \in K^\Gamma.$$

That is, the Γ orbit decomposes as a disjoint union of copies of $\mathbb{F}_2$ orbits.

The above theorem is due to Gaboriau-Lyons [GL09] for the case $(K, \kappa) = ([0,1], \texttt{Leb})$, and Bowen [Bow19] for the finitary cases. See also Houdayer [HOU11] for further useful background.

Question 3. Let G be a nonamenable group, and $\mathscr{P}$ the Poisson point process on G. Is there a free and ergodic action $\mathbb{F}_2 \curvearrowright (\mathbb{M}_0, \mathscr{P}_0)$ such that $\mathbb{F}_2\omega \subseteq [\omega]_\mathcal{R}$ for almost every ω?

Note that this implies one can construct a (directed, edge labelled) factor graph on the Poisson point process such that every connected component is a 4-regular tree.

One can ask the question for point processes other than the Poisson. In the case of IID labelled point processes this can be solved by directly applying work of Bowen, Hoff, and Iona [BHI18]. In that paper they solve the Measurable von Day problem for *Bernoulli extensions* of nonamenable pmp cbers over a probability space (K, κ). If Π is a point process, then the Palm equivalence relation of $(K, \kappa)^\Pi$ is exactly the Bernoulli extension of Π's Palm equivalence relation over (K, κ).

Remark 25. Ergodicity of the action in Question 3 translates to saying that the connected components of the associated factor graph are *indistinguishable* in the sense of Lyons and Schramm [LS99] (see also [Mar12] for further information).

1.7 Nonamenability

We now describe another characterisation of nonamenability of a group in terms of the associated unitary representation. We have no grand application of the observation, but simply record it for posterity.

Definition 27. Let $G \curvearrowright (X, \mu)$ be a measure preserving (mp) action. Its *Koopman representation* is the unitary representation π of G on $L^2(X, \mu)$ defined by

$$(\pi(g)f)(x) := f(g^{-1}x).$$

We simply write $L^2(X)$ if the measure μ is understood.

Let $L^2_0(X) = \{f \in L^2(X) \mid \int_X f(x)d\mu(x) = 0\}$ denote the G-invariant subspace of mean zero functions. Note that $L^2_0(X) = L^2(X)$ if the underlying measure μ has $\mu(X) = \infty$.

An *almost invariant sequence* in $L^2_0(X)$ is a sequence of unit vectors f_n such that

$$\lim_{n\to\infty} \|\pi(g)f_n - f_n\| = 0 \text{ for all } g \in G.$$

We say the the action $G \curvearrowright (X, \mu)$ *has spectral gap* if it has no almost invariant sequences.

[14]That is, κ is not the Dirac mass δ_k for some $k \in K$

For further details, see the survey paper of Bekka [Bek18].

Recall that G is amenable if and only if its regulat representation contains an almost invariant sequence.

Proposition 15. A group G is nonamenable if and only if the Poisson point process action $G \curvearrowright (\mathbb{M}, \mathscr{P})$ on it has spectral gap.

If G is discrete, then one should interpret the above statement as referring to the Bernoulli shift $G \curvearrowright (\{0,1\}^G, \mathtt{Ber}(p)^{\otimes G})$. In this case, the proposition is proved by expressing $L^2_0(\{0,1\}^G)$ as a direct sum of copies of the regular representation $\ell^2(G)$ and subregular representations. See Section 2.3.1 of Kerr and Li's book [KL16] for further details, and Lyons-Nazarov [LN11] for a particularly cool application of this fact.

In the nondiscrete case we appeal to an alternative decomposition of $L^2(\mathbb{M}, \mathscr{P})$ proved by Last and Penrose in [LP11].

If $\mathcal{H}$ is a Hilbert space over $\mathbb{R}$, we denote its nth tensor power by $\mathcal{H}^{\otimes n}$, with the convention that $\mathcal{H}^0 = \mathbb{R}$. We denote by $S^n(\mathcal{H})$ the subspace generated by the symmetric tensors.

The Koopman representation turns products of measure spaces into tensor products: that is, $L^2((X_1,\mu_1) \otimes (X_2,\mu_2)) = L^2(X_1,\mu_1) \otimes L^2(X_2,\mu_2)$. In the analogous identification for $L^2(G,\lambda)^{\otimes n}$, the symmetric tensors $S^n(L^2(G))$ are identified with the space of L^2 functions on G^n which are invariant under permutation of their variables.

Theorem 22 (Last-Penrose [LP11])**.** Let $\mathscr{P}$ denote the Poisson point process on G of unit intensity. Then the Koopman representation decomposes as

$$L^2(\mathbb{M}, \mathscr{P}) = \bigoplus_{n \geq 0} S^n(L^2(G)).$$

Remark 26. It should be stressed that Last and Penrose work with Poisson point processes in full generality on more-or-less arbitrary measure spaces, not merely the special case of lcsc groups with Haar measure. In particular, one also gets a similar decomposition of the Koopman representation of the IID Poisson on G.

The decomposition is achieved by specifying a rule that associates to functions $F : (\mathbb{M}, \mathscr{P}) \to \mathbb{R}$ a symmetric function $T_n F : G^n \to \mathbb{R}$.

Fix $g \in G$. We define the following *difference operator* D_g on measurable functions $F : \mathbb{M} \to \mathbb{R}$ in the following way:

$$(D_g F)(\omega) = F(\omega \cup \{g\}) - F(\omega).$$

Higher order difference operators $D^n_{g_1, g_2, \ldots, g_n}$ are defined inductively for $n \geq 2$ by

$$D^n_{g_1,g_2,\ldots,g_n} F = D^1_{g_1}(D^{n-1}_{g_2,\ldots,g_n} F).$$

Given $F \in L^2(\mathbb{M}, \mathscr{P})$, define $T_n F \in S^n(L^2, G)$ for $n \geq 1$ by

$$(T_n F)(g_1, g_2, \ldots, g_n) := \mathbb{E}_{\mathscr{P}}[D^n_{g_1,g_2,\ldots,g_n} F(\Pi),]$$

where Π denotes the unit intensity Poisson point process on G. Define $T_0 F = \mathbb{E}_{\mathscr{P}}[F(\Pi)]$.

With a great deal of work, it can be shown that the map $F \mapsto (T_0 F, T_1 F, \ldots)$ isometrically identifies $L^2(\mathbb{M}, \mathscr{P})$ with $\bigoplus_{n \geq 0} S^n(L^2(G))$ with its norms rescaled. Note that this map is also equivariant, so gives the desired decomposition as unitary representations.

Proof of proposition. Simply observe that $S^n(L^2(G))$ is a subrepresentation of $L^2(G)^{\otimes n}$ by definition, which is in turn a subrepresentation of $L^2(G)^{\oplus \mathbb{N}}$. Thus

$$L^2_0(\mathbb{M}, \mathscr{P}) \text{ is a subrepresentation of } L^2(G)^{\oplus \mathbb{N}}.$$

Now recall that a representation π has almost invariant vectors if and only if $\pi^{\oplus \mathbb{N}}$ does, finishing the proof.

□

Question 4. For those in the know, one sees that the Koopman representation of the Poisson point process on G is the same as the Koopman representation of the Gaussian action associated to the regular representation of G. In the discrete world, the Gaussian action associated to $\ell^2(\Gamma)$ is the Bernoulli shift $\Gamma \curvearrowright [0,1]^\Gamma$.

The quesion is: what is the relationship between the Gaussian action associated to $L^2(G)$ and the Poisson point process on G?

Chapter 2

Intermezzo: metric properties of $\mathbb{M}(X)$ and weak convergence

The following fact is the most basic requirement for a well-behaved probability theory:

Theorem 23 (See blah)**.** If X is a complete and separable metric space, then $\mathbb{M}(X)$ is a Borel subset of a complete and separable metric space $\mathcal{M}^{\#}(X)$, and is thus a standard Borel space.

Note that configurations $\omega \in X$ can be viewed as measures on X, by defining $\omega(A) = |\omega \cap A|$. So configurations form particular examples of *locally finite measures* on X, and $\mathcal{M}^{\#}(X)$ will be the space of such measures. In this language, a point process is a particular example of a *random measure*. Probabilists are interested in other examples of random measures[1], and have thus developed a framework suitable to handle all their cases of interest. There is no need to reinvent the wheel, so we adopt their framework with small notational adaptations.

We assume the reader is at least passingly familiar with weak converge of measures on metric spaces. Recall:

Definition 28. Let $\mathcal{M}(X)$ denote the space of *totally finite* measures η on X, that is, those with $\eta(X) < \infty$.

The *Prokhorov metric* d_{prok} on $\mathcal{M}(X)$ is

$$d_{\text{prok}}(\eta, \eta') = \inf\{\varepsilon \geq 0 \mid \text{for all Borel } A \subseteq X, \eta(A) \leq \eta'(A^\varepsilon) + \varepsilon \text{ and } \eta'(A) \leq \eta(A^\varepsilon) + \varepsilon\},$$

where A^ε is the ε*-halo* of A, that is,

$$A^\varepsilon = \{x \in X \mid d(x, A) < \varepsilon\}.$$

If η is a totally finite measure on X, then a η*-continuity set* is a subset $A \subseteq X$ with the property that $\eta(\partial A) = 0$, where ∂A denotes the *topological boundary* of A.

A sequence of totally finite measures η_n *weakly converges* to η if either of the following conditions hold:

WC1 for all continuous and bounded functions $f : X \to \mathbb{R}$,

$$\lim_{n\to\infty} \int_X f(x) d\eta_n(x) = \int_X f(x) d\eta(x).$$

[1] and, for that matter, non-invariant point processes

WC2 for all η-continuity sets $A \subseteq X$,

$$\lim_{n\to\infty} \eta_n(A) = \eta(A).$$

Remark 27. The Prokhorov metric metrises this convergence notion (that is, η_n converges weakly to η if and only if $d_{\text{prok}}(\eta_n, \eta)$ converges to zero).

The equivalence of **WC1** and **WC2** is usually referred to as *the Portmanteau theorem.*

The definition involving continuity sets will have preeminence for us. To explain the name: note that the indicator function $\mathbb{1}_A : X \to \{0,1\}$ is continuous η almost everywhere if and only if A is an η-continuity set.

We often make use of the following well-known fact:

Lemma 4. If $\Phi : X \to Y$ is a continuous map of metric spaces, then Φ preserves weak limits: if μ_n is a sequence of Borel probability measures on X weakly converging to μ, then $\Phi_*\mu_n$ weakly converges to $\Phi_*\mu$.

Moreover, the same is true if Φ is merely continuous μ almost everywhere.

Proof. The first statement is immediate: if $f : Y \to \mathbb{R}$ is a continuous and bounded function, then $f \circ \Phi : X \to \mathbb{R}$ is continuous and bounded as well, so

$$\lim_{n\to\infty} \int_Y f d\Phi_*\mu_n = \lim_{n\to\infty} \int_X f \circ \Phi d\mu_n = \int_X f \circ \Phi d\mu = \int_Y f d\Phi_*\mu.$$

For the second statement we work with the definition involving continuity sets. Let $A \subseteq Y$ be a $\Phi_*\mu$ continuity set, that is, assume $\mu(\Phi^{-1}(\partial A)) = 0$. Let $D_\Phi = \{x \in X \mid \Phi \text{ is discontinuous at } x\}$ denote the discontinuity set of Φ. One can show that for any $A \subseteq Y$ that $\partial\Phi^{-1}(A) \subseteq \Phi^{-1}(\partial A) \cup D_\Phi$, and so

$$\mu(\partial\Phi^{-1}(A)) \leq \mu(\Phi^{-1}(\partial A)) + \mu(D_\Phi) = 0,$$

that is, $\Phi^{-1}(A)$ is a μ-continuity set. Therefore

$$\lim_{n\to\infty} \Phi_*\mu_n(A) = \lim_{n\to\infty} \mu_n(\Phi^{-1}(A)) = \mu(\Phi^{-1}(A)) = \Phi_*\mu(A),$$

as desired. □

Definition 29. Let $\mathcal{M}^\#(X)$ denote the space of *boundedly finite* measures, that is, those Borel measures η on X that are finite on metrically bounded subsets of X.

Fix a basepoint $x_0 \in X$. Let

$$\eta^{(r)}(A) := \eta(A \cap B(x_0; r))$$

denote the restriction of a boundedly finite measure η to the r-ball about x_0. Note that $\eta^{(r)}$ is therefore an element of $\mathcal{M}(X)$.

We now define a metric $d^\#$ on $\mathcal{M}^\#(X)$:

$$d^\#(\eta, \eta') = \int_0^\infty e^{-r} \frac{d_{\text{prok}}(\eta^{(r)}, \eta'^{(r)})}{1 + d_{\text{prok}}(\eta^{(r)}, \eta'^{(r)})} dr.$$

A sequence of boundedly finite measures η_n *weak-# converges* to η if any of the following conditions hold:

WHC1 for all continuous and bounded functions $f : X \to \mathbb{R}$ which vanish outside a bounded set,

$$\lim_{n\to\infty} \int_X f(x) d\eta_n(x) = \int_X f(x) d\eta(x).$$

WHC2 for all bounded η-continuity sets $A \subseteq X$,

$$\lim_{n\to\infty} \eta_n(A) = \eta(A).$$

WHC3 there exists a sequence r_k of radii increasing to infinity such that for every $k \in \mathbb{N}$

$$\eta_n^{(r_k)} \text{ converges weakly to } \eta^{(r_k)}.$$

Remark 28. The space we've defined is obviously *extremely* metrically dependent (recall that every metric is topologically equivalent to a bounded metric). However, our case of interest is proper left-invariant metrics on locally compact groups, which are all coarsely equivalent and thus have a well-defined notion of metrically boundedness.

In case X is locally compact, then weak-# convergence is equivalent to vague convergence.

Theorem 24. The space $\mathcal{M}^{\#}(X)$ equipped with the $d^{\#}$ metric is complete and separable. Its Borel structure is exactly such that the mass measuring[2] functions $N_A : X \to \mathbb{N}_0 \cup \infty$ given by $\eta \mapsto \eta(A)$ are measurable, where A is an arbitrary Borel subset of X.

Remark 29. The Borel structure on $\mathcal{M}^{\#}(X)$ can be generated by an even smaller collection of mass measuring functions: one only needs to look at N_A where A ranges over a semiring of bounded Borel sets that generate the Borel structure on X.

We will require the following more explicit explanation of what weak-# convergence is:

Definition 30. Let $\omega \in \mathbb{M}(X)$ be a configuration. We call another configuration $\omega' \in \mathbb{M}(X)$ a (ε, R)-*wobble* of ω (where $\varepsilon, R > 0$ are some parameters) if $\omega^{(R)}$ is in bijection with $\omega'^{(R)}$, and moreover this bijection $\sigma : \omega^{(R)} \to \omega'^{(R)}$ can be chosen in such a way that $d(x, \sigma(x)) < \varepsilon$ for all $x \in \omega^{(R)}$.

One direction of the following lemma is immediate, the converse is less elementary and can be found in [VJ03] as Proposition A2.6.II:

Lemma 5. A sequence of configurations ω_n converges to ω with respect to $d^{\#}$ *if and only if* there are sequences $\varepsilon_n \to 0$ and $R_n \to \infty$ such that each ω_n is a (ε_n, R_n)-wobble of ω.

We can now discuss *weak convergence* of point processes. View $\mathbb{M}(X)$ as a subset of $\mathcal{M}^{\#}(X)$ equipped with the $d^{\#}$ metric, and recall that a point process is a probability measure on $\mathbb{M}(X)$. This is what we mean by a sequence of point processes weakly converging.

Note that the weak limit of a sequence of point processes μ_n will (a priori) be a probability measure on $\mathcal{M}^{\#}(X)$, *not* on $\mathbb{M}(X)$. That is, a point process might converge to a random measure which is not a point process. It's easy to see that the only thing that can go wrong is mass accumulation: the limit measure will be a random atomic measure, but some atoms might have mass larger than one.

[2] earlier we called these "point counting" functions, because that's a more suitable name when the measure is atomic

Definition 31. A *counting measure* on X is a measure η with $\eta(A) \in \mathbb{N}_0$ for all bounded Borel subsets $A \subseteq X$. A *simple* counting measure is a measure η with $\eta(\{x\}) = 0$ or 1 for all $x \in X$.

If η is a counting measure, then its *support* is $\text{support}(\eta) = \{x \in X \mid \eta(\{x\}) > 0\}$. That is, support($\eta$) is η with the multiplicities removed.

Example 22. Let $\{X_k\}_{k\in\mathbb{Z}}$ denote an IID sequence of uniform $[0, 1]$ random variables. Consider the following sequence of point processes:

$$\Pi_n = \mathbb{Z} \cup \{k + \frac{X_k}{n}\}.$$

In words: take two copies of $\mathbb{Z}$, where you wobble all the points of one copy by smaller and smaller amounts (this is not a point process proper in our sense, as it is not invariant, but one can take a uniform $[0, 1]$ shift of Π_n if one insists).

Then Π_n weakly converges to the deterministic measure μ given by $\mu(A) = 2|A \cap \mathbb{Z}|$.

Remark 30. In this language, what we've been calling point processes are *random simple counting measures*, and the comment above states that the weak limit of random simple counting measures, if it exists, will be a possibly non-simple random counting measure.

In the literature one sometimes sees random counting measures referred to as "point processes", and random simple counting measures as "simple point processes".

Definition 32. Let (X, d) be a csms, and (μ_n) a sequence of Borel probability measures on X. The sequence is *uniformly tight* if for every $\varepsilon > 0$ there exists a compact set $K \subseteq X$ such that $\mu_n(X \setminus K) < \varepsilon$ for all $n \in \mathbb{N}$.

Recall from Prokhorov's theorem that a sequence (μ_n) is uniformly tight if and only if its closure $\overline{(\mu_n)}$ is compact.

There is a more explicit form of uniform tightness that we will use repeatedly:

Theorem 25 (cite)**.** Suppose X is a *locally compact*[3] csms. A sequence of point processes (Π_n) on X is uniformly tight if and only if for every closed ball $B \subseteq X$ and any $\varepsilon > 0$ there exists an $M > 0$ such that

$$\mathbb{P}[N_B \Pi_n > M] < \varepsilon \text{ for all } n \in \mathbb{N}.$$

Lemma 6. Let

$$H_\delta = \{\omega \in \mathbb{M}(X) \mid d(x, y) \geq \delta \text{ for all distinct } x, y \in \omega\}$$

denote the space of *δ-uniformly-separated* configurations. Then H_δ is compact in $\mathbb{M}(X)$.

Probabilists seem to use the term "hard-core" for configurations with this property hence our choice of letter. I am not aware of a standard notation for this space.

The previous lemma is proved using the following basic fact:

Lemma 7. Let (X, d) denote a compact metric space. Then for all $\delta > 0$ there exists some $C > 0$ such that $|A| \leq C$ for any δ-separated subset $A \subseteq X$.

[3]There is a slightly more complicated version of the theorem for general Polish spaces, but we will not use it, so do not state it

The above discussion has been rather abstract. We now outline an equivalent interpretation of weak convergence that will be much more useful in certain applications.

Definition 33. Let Π be a point process with law μ. A *stochastic continuity set* of Π is a Borel subset $V \subseteq G$ of the ambient space such that $\mathbb{P}[\Pi \cap \partial V \neq \varnothing] = 0$. Equivalently, it is a subset such that its point counting function $N_V : \mathbb{M} \to \mathbb{N}_0 \cup \{\infty\}$ is continuous μ almost everywhere.

Let $\boldsymbol{V} = (V_1, V_2, \ldots, V_k)$ denote a collection of stochastic continuity sets for Π.

The *finite dimensional distributions* of Π are the random vectors

$$N_{\boldsymbol{V}}(\Pi) = (N_{V_1}\Pi, N_{V_2}\Pi, \ldots, N_{V_k}\Pi),$$

where $\boldsymbol{V}$ runs over all possible collections of stochastic continuity sets.

Remark 31. The sets

$$\{\omega \in \mathbb{M} \mid N_{\boldsymbol{V}}(\omega) = \boldsymbol{\alpha}\}, \text{ where } \boldsymbol{\alpha} = (\alpha_1, \alpha_2, \ldots, \alpha_k) \in \mathbb{N}_0^k$$

should be thought of as analogous to the *cylinder sets* in the space $\{0,1\}^\Gamma$, where Γ is a discrete group.

Note that if $\boldsymbol{V}$ is a family of stochastic continuity sets for Π, then $N_{\boldsymbol{V}}$ is continuous μ almost everywhere. Thus by the earlier fact on weak limits and continuous functions, we see that weak convergence of point processes implies weak convergence of the finite dimensional distributions. The surprising fact is that the converse is true:

Theorem 26 (See Theorem 11.1.VII of [DVJ07])**.** A sequence Π_n of point processes weakly converges to Π *if and only if* for all collections $\boldsymbol{V} = (V_1, V_2, \ldots, V_k)$ of stochastic continuity sets of Π, the finite dimensional distributions $N_{\boldsymbol{V}}(\Pi_n)$ weakly converge to $N_{\boldsymbol{V}}(\Pi)$.

For this to make any sense at all, it must be the case that the finite dimensional distributions *determine* a point process. That is, if two point processes Π and Π' have $N_{\boldsymbol{V}}(\Pi) \overset{d}{=} N_{\boldsymbol{V}}(\Pi')$ for all collections of stochastic continuity sets $\boldsymbol{V}$, then $\Pi \overset{d}{=} \Pi'$. This is proved using the following lemma, which states that there is an abundance of continuity sets:

Lemma 8. Let Π be a point process with law μ. Then

- for all $g \in G$, there are at most countably many $r > 0$ such that the open ball $B_G(g, r)$ is *not* a stochastic continuity set of Π, and
- for all $\omega \in \mathbb{M}$, there are at most countably many $r > 0$ such that the open ball $B_{\mathbb{M}}(\omega, r)$ is *not* a μ continuity set.

In particular, both G and $\mathbb{M}$ admit *topological bases* consisting of μ stochastic continuity sets / μ continuity sets (respectively).

Proof. The method is the same in both cases, so we only write the proof for the first statement. The idea of the proof is that that there cannot be so many stochastic continuity sets, else local finiteness will be contradicted. It is enough to prove that for every $r > 0$ and $\varepsilon > 0$ there exists only finitely many $r_1, r_2, \ldots$ in $(0, r)$ such that

$$\mathbb{P}[\Pi \cap \partial B(0, r_i) \neq \varnothing] > \varepsilon \text{ for all } i.$$

Suppose not. That is, suppose we have $r, \varepsilon > 0$ and infinitely many $\{r_n\} \subset (0, r)$ satisfying the above equation. Then

$$\varepsilon \leq \limsup_{n\to\infty} \mathbb{P}[\Pi \cap \partial B(0, r_n) \neq \varnothing] \leq \mathbb{P}\left[\limsup_{n\to\infty} (\Pi \cap \partial B(0, r_n) \neq \varnothing)\right].$$

Recall that the lim sup of a sequence of events is the event that they occur infinitely often. So we see

$$\{\Pi \cap \partial B(0, r_n) \neq \varnothing \text{ for infinitely many } n\} \subseteq \{|\Pi \cap B(0, r)| = \infty\}.$$

We've shown that with positive probability, Π has infinitely many points in $B(0, r)$, a contradiction by local finiteness. □

Note that the continuity sets form an algebra, and the cylinder sets $\{\omega \in \mathbb{M} \mid N_{\boldsymbol{V}}(\omega) = \boldsymbol{\alpha}\}$ are continuity sets when $\boldsymbol{V}$ is a collection of stochastic continuity sets. As a measure is determined by its values on an algebra that generates the Borel sigma algebra, we therefore see that point processes are determined by their finite dimensional distributions. With a bit more work (see [VJ03] Proposition A2.3.IV and [DVJ07] Corollary 11.1.III, Theorem 11.I.VII), one can prove:

Lemma 9. Let Π be a point process. Then there exists a *countable* family $\{V_i\}_{i\in\mathbb{N}}$ of *metrically bounded* and disjoint Borel subsets $V_i \subseteq G$ such that Π_n weakly converges to Π *if and only if* $N_{\boldsymbol{V}}\Pi_n$ weakly converges to $N_{\boldsymbol{V}}\Pi$ were $\boldsymbol{V}$ ranges over all finite subcollections of $\{V_i\}$.

In particular, weak convergence can be verified by a *countable* collection of statements, each of which only requires one to observe the process in compact windows.

The following lemma is a simpler case of exercise 13.2.2 in [DVJ07], and is presumably known with a more elegant proof. The technique will be used for a later proof, so we include it.

Proposition 16. Suppose μ^n is a sequence of point processes that weakly converges to a finite intensity process μ. Then μ^n has finite intensity for all but finitely many n, moreover, the Palm measures μ_0^n weakly converge to μ_0.

Proof. Let $A \subseteq \mathbb{M}_0$ be a μ-continuity set, and $U \subseteq G$ a stochastic μ continuity set of unit volume. Recall that

$$\mu_0(A) = \frac{1}{\text{intensity}\,\mu}\mathbb{E}_\mu\left[\#\{g \in U \mid g^{-1}\omega \in A\}\right]$$

Claim. For every $k \in \mathbb{N}$, the function $\omega \mapsto \max\{\#\{g \in U \mid g^{-1}\omega \in A\}, k\}$ is continuous μ almost everywhere.

This function can only be discontinuous on the boundary of

$$A^{(l)} = \{\omega \in \mathbb{M} \mid \#\{g \in U \mid g^{-1}\omega \in A\} = l\}.$$

Note that the claim applied with $A = \mathbb{M}_0$ together with the monotone convergence theorem proves in particular that

$$\lim_{n\to\infty} \text{intensity}\,\mu_n = \text{intensity}\,\mu,$$

using the definition of weak convergence involving integrals of continuous bounded functions. The same sort of argument proves the proposition itself.

Observe that

$\mu_0(\partial A) = 0$	By assumption, so
$\mu_0([\partial A]) = 0$	as saturations of null sets are null in a countable groupoid,
$\mu(G\partial A) = 0$	by Proposition 5.

We show that $\partial A^{(l)} \cap \{N_{\partial U}\omega = 0\} \subseteq G\partial A$ for all $l \geq 1$, establishing the claim.

Suppose $\omega \in \partial A^{(l)} \cap \{N_{\partial U}\omega = 0\}$. Then we can find two sequences ω_n, ω'_n both converging to ω such that $\omega_n \in A^{(l)}$ and $\omega'_n \notin A^{(l)}$ for every $n \in \mathbb{N}$. We take these to be (ε_n, R_n)-wobbles of ω.

We see that (for large n) the configurations $\omega, \omega_n, \omega'_n$ are all approximately equal inside U. See Figure 2.

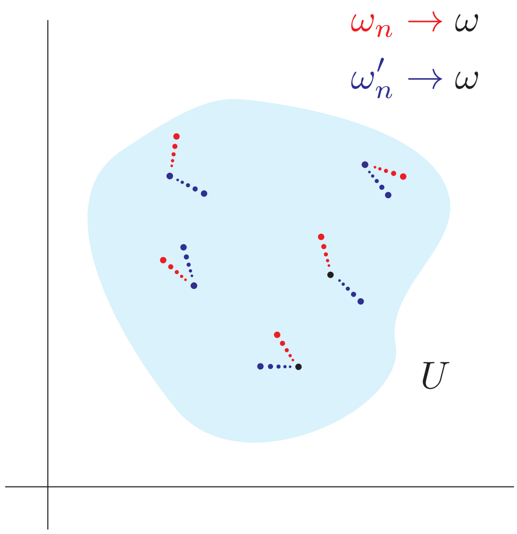

We will refer to the points $g \in U \cap \omega$ such that $g^{-1}\omega \in A$ as *A-points of ω* and likewise for ω_n and ω'_n.

Now for every (large) n, the number of A points of ω_n in U is l, and the number of A points of ω'_n in U is some (bounded) number other than l. Since the configurations are a small wobble of ω then, we can find $g_n \in \omega \cap U$ such that the corresponding points x_n of ω_n and y_n of ω'_n are an A point and *not* an A point, respectively.

As g_n ranges over a finite set $\omega \cap U$, we can choose $g \in \omega \cap U$ and a subsequence (n_k) such that $g_{n_k} = g$ for every $k \in \mathbb{N}$. Then

$$x_{n_k}^{-1}\omega_{n_k} \to g^{-1}\omega, \text{ and } y_{n_k}^{-1}\omega'_{n_k} \to g^{-1}\omega,$$

which witnesses that $g^{-1}\omega \in \partial A$, so $\omega \in G\partial A$, as desired. □

Chapter 3

Computing cost in the weak limit

3.1 Weak factoring and Abert-Weiss for point processes

We have seen that cost is monotone under factor maps. We introduce the following weakening of factoring, inspired by the concept of weak containment of pmp actions:

Definition 34. Let Π and Υ be point processes. Then Π *weakly factors* onto Υ if there is a sequence Φ^n of factors of Π such that $\Phi^n(\Pi)$ weakly converges to Υ.

Question 5. Is weak factoring transitive? That is, if Π weakly factors onto Υ, then does Π weakly factor onto any weak factor of Υ?

We are able only to prove the following limited case of transitivity. Fortunately it is enough to deduce our desired statements about cost, but unfortunately it adds some tedious complications to their proofs.

Proposition 17. Suppose Π and Υ are point processes, and Π weakly factors onto Υ. Then Π weakly factors onto any *thinning* of Υ too.

Proof. Let $\Phi^n(\Pi)$ weakly converge to Υ, and θ be a thinning of Υ, determined by a subset A of $(\mathbb{M}_0, \nu_0)$.

Let θ^A be a thinning of Υ, determined by a subset $A \subseteq (\mathbb{M}_0, \nu_0)$ as in Proposition 6.

Let $A_n \subseteq \mathbb{M}_0$ be ν_0-continuity sets such that

$$\nu_0(A \triangle A_n) < \frac{1}{\lambda(B(0,n)2^n}.$$

Claim. The thinnings $\theta^{A_n} : (\mathbb{M}, \nu) \to \mathbb{M}$ are continuous ν almost everywhere, and converge pointwise to θ^A.

With the claim in hand, the proof is finished as follows: almost sure convergence implies weak convergence, so $\theta^{A_n}(\Upsilon)$ converges weakly to $\theta^A(\Upsilon)$. By Lemma 4, for every $n \in \mathbb{N}$ we have $\theta^{A_n}(\Phi^k(\Pi))$ converging weakly to $\theta^{A_n}(\Upsilon)$ as $k \to \infty$. By choosing a suitable subsequence of k's and n's we see that Π weakly factors onto $\theta(\Upsilon)$.

For the claim, note that $\nu(G\partial A_n) = 0$, and θ^{A_n} is continuous off that set. For pointwise convergence, note that

$$\begin{aligned}
\mathbb{P}[\exists x \in B(0,n) \text{ such that } x^{-1}\Upsilon \in A\triangle A_n] &\leq \mathbb{E}[\#\{x \in B(0,n) \cap \theta^{A\triangle A_n}(\Upsilon)\}] \\
&= \text{intensity}(\theta^{A\triangle A_n}(\Upsilon))\lambda(B(0,n)) \\
&\leq \frac{1}{\lambda(B(0,n)2^n}\lambda(B(0,n) \\
&= \frac{1}{2^n}.
\end{aligned}$$

So Borel-Cantelli on these events implies pointwise convergence. ☐

Recall that an arbitrary point process factor decomposes as a thinning and a thickening, so to prove the general statement it suffices to prove that if Π weakly factors onto Υ, then Π weakly factors onto any thickening of Υ.

Lemma 10. Suppose Π_n weakly converges to Π. Then $[0,1]^{\Pi_n}$ weakly converges to $[0,1]^{\Pi}$.

Proof. This can be seen, for instance, by verifying that the finite dimensional distributions of $[0,1]^{\Pi_n}$ weakly converge to those of $[0,1]^{\Pi}$.

Recall that a $[0,1]$-marked point process on G is just a particular kind of point process on $G \times [0,1$. It therefore suffices to check weak convergence of the finite dimensional distributions against stochastic continuity sets of $[0,1]^{\Pi}$ in product form (see [citation]).

To that end, let $\boldsymbol{V} = (V_1, V_2, \ldots, V_k)$ denote a collection of stochastic continuity sets for Π, and $[0,\boldsymbol{p}) = ([0,p_1),[0,p_2),\ldots,[0,p_k))$ a family of intervals in $[0,1]$. We denote by $\boldsymbol{V} \times [0,\boldsymbol{p}) = (V_1 \times [0,p_1),\ldots,V_k \times [0,p_k))$ the stochastic continuity set of $[0,1]^{\Pi}$. Fix an integral vector $\boldsymbol{\alpha} \in \mathbb{N}_0^k$. We must show that $\mathbb{P}[N_{\boldsymbol{V}\times[0,\boldsymbol{p})}[0,1]^{\Pi_n} = \boldsymbol{\alpha}]$ converges to $\mathbb{P}[N_{\boldsymbol{V}\times[0,\boldsymbol{p})}[0,1]^{\Pi} = \boldsymbol{\alpha}]$.

We find the following explicit expression simply by conditioning on $\boldsymbol{\beta}$, the total number of points appearing in $\boldsymbol{V}$:

$$\begin{aligned}
\mathbb{P}[N_{\boldsymbol{V}\times[0,\boldsymbol{p})}[0,1]^{\Pi_n} = \boldsymbol{\alpha}] &= \sum_{\boldsymbol{\beta}\geq\boldsymbol{\alpha}} \mathbb{P}[N_{\boldsymbol{V}\times[0,\boldsymbol{p})}[0,1]^{\Pi_n} \mid N_{\boldsymbol{V}}\Pi_n = \boldsymbol{\beta}]\mathbb{P}[N_{\boldsymbol{V}}\Pi_n = \boldsymbol{\beta}] \\
&= \sum_{\boldsymbol{\beta}\geq\boldsymbol{\alpha}} \prod_{i=1}^{k} p_i^{\alpha_i}(1-p_i)^{\beta_i-\alpha_i}\mathbb{P}[N_{\boldsymbol{V}}\Pi_n = \boldsymbol{\beta}],
\end{aligned}$$

where by $\boldsymbol{\beta} \geq \boldsymbol{\alpha}$ we mean that $\beta_i \geq \alpha_i$ for each entry.

There is an identical expression for Π (simply replace all instances of Π_n by Π). The conclusion follows, as $\mathbb{P}[N_{\boldsymbol{V}}\Pi_n = \boldsymbol{\beta}]$ converges to $\mathbb{P}[N_{\boldsymbol{V}}\Pi = \boldsymbol{\beta}]$ for all $\boldsymbol{\beta}$. ☐

Theorem 27. Let Π and Υ be *any* point processes on an amenable group G. Then $[0,1]^{\Pi}$ weakly factors onto Υ.

Proof. We can assume that Π is a random net, as follows: certainly $[0,1]^{\Pi}$ factors onto some net Π' by Proposition 10, and therefore onto $[0,1]^{\Pi'}$ by the replication trick. Then if $[0,1]^{\Pi'}$ weakly factors onto Υ, so too does $[0,1]^{\Pi}$.

The random net $[0,1]^{\Pi}$ admits a one-ended clumping, which we can use to construct a sequence of coarser and coarser partitions $\mathcal{P}_n$ on G as in Section 1.5. In each cell of

$\mathcal{P}_n$, we select a point randomly, and use it to place down an independent sample from Υ. Now if W is any bounded window, it meets only finitely many of the Voronoi cells of Pi. Hence for large n, it will be *entirely* contained in a single cell of the partition. This occurs with high probability, and on this event the finite dimensional distributions of the factor coincide with Υ's, proving weak convergence.

□

The following statement is due to Abért and Weiss [AW13] for discrete groups, we extend it to point processes:

Theorem 28. Let Π be an essentially free point process. Then it weakly factors onto $[0,1]^\Pi$, its own IID.

Proof. It suffices to show that Π weakly factors onto $[d]^\Pi$, where $[d] = \{1, 2, \dots, d\}$ is equipped with the uniform measure, as $[d]^\Pi$ weakly converges to $[0,1]^\Pi$ as $d \to \infty$. We will do this by constructing factor $[d]$-*labellings* $\mathscr{C}_n$ of Π such that $\mathscr{C}_n(\Pi)$ weakly converges to $[d]^\Pi$.

To do this, we'll use the second moment method, hewing close to the original Abert-Weiss recipe.

The strategy will be as follows. Consider the set of $[d]$-labellings of Π. We will consider a model that produces a *random element* of this space. We will show that it satisfies certain constraints with positive probability. In particular, there must exist a $[d]$-labelling satisfying those constraints. By adjusting the parameters of this model, one can produce the desired sequence $\mathscr{C}_n$.

Fix a *countable* weak convergence determining family $\{V_i\}$ as discussed at Lemma 9, so that the sets $V_i \subset G \times [d]$ are bounded stochastic continuity sets for $[d]^\Pi$. We will construct a sequence of factor colourings $\mathscr{C}_n$ of Π such that for fixed k,

$$N_{\boldsymbol{V}_k}(\mathscr{C}_n\Pi) \text{ converges weakly to } N_{\boldsymbol{V}_k}([d]^\Pi),$$

where $\boldsymbol{V}_k = (V_1, V_2, \dots, V_k)$.

Set $W_k = \bigcup_{i \leq k} V_k$ to be the total window. Formally this is a subset of $G \times [d]$, but we view it as a subset of G. For $\varepsilon > 0$ arbitrary, we choose $\delta > 0$ so small that the following properties are true, where μ denotes the law of Π:

$$\mu(\{\omega \in \mathbb{M} \mid \text{for all } g, h \in \omega \cap W_k, g \neq h \text{ implies } d(g^{-1}\omega, h^{-1}\omega) > \delta\}) > 1 - \varepsilon$$

and

$$(\mu \otimes \mu)(\{(\omega, \omega') \in \mathbb{M} \times \mathbb{M} \mid \text{for all } (g,h) \in (\omega \cap W_k) \times (\omega' \cap W_k), \text{ we have } d(g^{-1}\omega, h^{-1}\omega) > \delta\}) > 1 - \varepsilon.$$

This is possible by *essential freeness* of Π: the sets in question increase as δ tends to zero to a set of full measure.

We now construct a *random* colouring $\mathscr{C}$ of Π in the following way: let

$$\mathbb{M}_0 = \bigsqcup_i D_i, \text{ where } \operatorname{diam}(D_i) < \delta.$$

be a partition of $\mathbb{M}_0$ into small measurable sets. By the correspondences we've described, any $[d]$-colouring of the sets D_i corresponds to a factor colouring $\mathscr{C} : \mathbb{M} \to [d]^{\mathbb{M}}$ in the following way:

$$\mathscr{C}(\omega) = \{(g, c) \in \omega \times [d] \mid g^{-1}\omega \in \mathbb{M}_0 \text{ is coloured by } c\}.$$

We look at such $\mathscr{C}$ when the D_i sets are coloured *uniformly at random* by elements of $[d]$. To emphasise: we are considering a *distribution* on *deterministic* colourings.

For an integral vector $\boldsymbol{\alpha} = (\alpha_1, \alpha_2, \ldots, \alpha_k) \in \mathbb{N}_0^k$, we set

$$T_\alpha = \{\omega \in [d]^{\mathbb{M}} \mid (N_{V_1}(\omega), \ldots, N_{V_k}(\omega)) = \alpha\}$$

to be the set of configurations whose point/colour statistics in W_k are prescribed by α.

Note that $\mathscr{C}_*\mu(T_\alpha)$ is a random variable (whose source of randomness is $\mathscr{C}$).

We use the second moment method to prove the existence of $\mathscr{C}$ such that for all $\alpha \in \mathbb{N}_0^k$ with $\|\alpha\|_\infty \leq M$,

$$|\mathscr{C}_*(\mu)(T_\alpha) - [d]^\mu(T_\alpha)| < \varepsilon.$$

Then any sequence of such colourings with k, M tending to infinity will witness that Π weakly factors onto $[d]^\Pi$.

Exchanging order of integration allows us to express the mean of $\mathscr{C}_*(\mu)(T_\alpha)$ as

$$\begin{aligned}\mathbb{E}\left[\mathscr{C}_*(\mu)(T_\alpha)\right] &= \mathbb{E}[\mu(\mathscr{C}^{-1}(T_\alpha))] \\ &= \mathbb{E}\left[\int_{\mathbb{M}} \mathbb{1}[\mathscr{C}(\omega) \in T_\alpha\right] d\mu(\omega) \\ &= \int_{\mathbb{M}} \mathbb{E}\left[\mathbb{1}[\mathscr{C}(\omega) \in T_\alpha]\right] d\mu(\omega).\end{aligned}$$

Note that for $\omega \in A_\delta$, all pairs of distinct points $g, h \in \omega$ from the window W_k have the property that $g^{-1}\omega$ and $h^{-1}\omega$ fall into different D_i sets, and are therefore assigned *independent* colours. Thus

$$\text{for } \omega \in A_\delta, \mathbb{E}\left[\mathbb{1}[\mathscr{C}(\omega) \in T_\alpha]\right] = [d]^\mu(T_\alpha).$$

As $\mu(A_\delta) > 1 - \varepsilon$, it follows that

$$|\mathbb{E}\left[\mathscr{C}_*\mu(T_\alpha)\right] - [d]^\mu(T_\alpha)| < 2\varepsilon.$$

We now work on the variance. Again, exchanging order of integration in a similar way to before allows us to express the mean of $(\mathscr{C}_*(\mu)(T_\alpha))^2$ as

$$\mathbb{E}\left[(\mathscr{C}_*(\mu)(T_\alpha))^2\right] = \iint_{\mathbb{M}\times\mathbb{M}} \mathbb{E}\left[\mathbb{1}[\mathscr{C}(\omega) \in T_\alpha \mathbb{1}[\mathscr{C}(\omega') \in T_\alpha\right] d\mu(\omega)\, d\mu(\omega').$$

By similar reasoning to before, for $(\omega, \omega') \in (A_\delta \times A_\delta) \cap B_\delta$, the colours one will see at points in W_k will be independent. Thus for such (ω, ω') we have

$$\mathbb{E}\left[(\mathscr{C}_*(\mu)(T_\alpha))^2\right] = \left([d]^\mu(T_\alpha)\right)^2$$

Note that $(A_\delta \times A_\delta) \cap B_\delta = (A_\delta \times \mathbb{M}_0) \cap (\mathbb{M}_0 \times A_\delta) \cap B_\delta$, so by the union bound $(\mu \otimes \mu)((A_\delta \times A_\delta) \cap B_\delta) > 1 - 3\varepsilon$.

Putting this together,

$$\operatorname{Var}(\mathscr{C}_*(\mu)(T_\alpha)) = \mathbb{E}\left[(\mathscr{C}_*(\mu)(T_\alpha))^2\right] - \left(\mathbb{E}\left[\mathscr{C}_*\mu(T_\alpha)\right]\right)^2 < 12\varepsilon.$$

We now apply Chebyshev's inequality which states that for any $c > 0$,

$$\mathbb{P}\left[|\mathscr{C}_*(\mu)(T_\alpha) - \mathbb{E}[\mathscr{C}_*(\mu)(T_\alpha)]| > c\right] < \frac{\operatorname{Var}(\mathscr{C}_*(\mu)(T_\alpha))}{c^2}.$$

Our bounds on the mean and the variance of $\mathscr{C}_*(\mu)(T_\alpha)$ and the choice $c = \varepsilon^{\frac{1}{3}}$ yield

$$\mathbb{P}\left[|\mathscr{C}_*(\mu)(T_\alpha) - [d]^\mu(T_\alpha)| > \varepsilon^{\frac{1}{3}} - 2\varepsilon\right] < 12\varepsilon^{\frac{1}{3}}.$$

Let E_α denote the event $\{|\mathscr{C}_*(\mu)(T_\alpha) - [d]^\mu(T_\alpha)| < \varepsilon\}$. Then by the union bound

$$\mathbb{P}[\bigcap_{\substack{\alpha\in\mathbb{N}_0^k \\ \|\alpha\|_\infty \leq M}} E_\alpha] \geq 1 - M^k\varepsilon^{\frac{1}{3}}.$$

In particular, by choosing ε sufficiently small, such a colouring exists. □

Remark 32. If weak factoring is transitive, then this can be combined with Theorem 27 to deduce:

Let Π and Υ be point processes on an amenable group G, with Π free. Then Π weakly factors onto Υ.

3.1.1 Factoring vs. IID labels

We have seen that all free point processes are able to *weakly* factor onto their own IID. It is natural to ask if all this hassle was worth it – can a point process always factor directly onto its own IID? The following observation was developed as part of a collection of works on various *thinning* and *thickening* questions suggests not:

Theorem 29 (Holroyd, Lyons, Soo [HLS11])**.** The Poisson point process cannot be split into two *independent* Poisson point processes of lower intensity without additional randomness.

More precisely, there does not exist a *deterministic* two colouring $\mathscr{C} : (\mathbb{M}, \mathscr{P}) \to \{0,1\}^\mathbb{M}$ such that $\mathscr{C}_*\mathscr{P}$ is the IID $\texttt{Ber}(p)$ labelled Poisson point process for $0 < p < 1$.

Example 23. Some point processes *can* factor onto their own IID. Note that taking the IID of a point process is idempotent, in the sense that

$$[0,1]^{[0,1]^\Pi} \cong ([0,1]^2)^\Pi \cong [0,1]^\Pi.$$

For an unlabelled example, one can simply *spatially implement* $[0,1]^\Pi$. That is, using the method sketched at Proposition 21 one can find an unlabelled point process Υ (abstractly) isomorphic to $[0,1]^\Pi$, and thus $[0,1]^\Upsilon \cong \Upsilon$.

3.2 Cost monotonicity for (certain) weak factors

In this section we will always assume G is compactly generated by $S \subset G$.

Question 6. Suppose Π weakly factors onto Υ. Is it true that $\text{cost}(\Pi) \leq \text{cost}(\Upsilon)$? That is, is cost monotone for weak factors?

This is the *real* theorem that we would like to prove. We are able only to prove the following theorem, which implies that cost is monotone for certain weak factors:

Theorem 30. Suppose Π^n is a sequence of point processes that weakly converge to Π. Then

$$\limsup_{n\to\infty} \operatorname{cost}(\Pi^n) \leq \operatorname{cost}(\Pi)$$

holds in the following cases:

1. If there exists $\delta, R > 0$ such that Π_n and Π are all δ uniformly separated and R coarsely dense.

2. If all the Π_n are free and Π is δ uniformly separated.

Moreover, the same statement is true if the point processes have labels from a *compact* mark space Ξ.

We will need an auxiliary lemma, which we will use again later:

Lemma 11. Let Π be a point process with law μ. Then for all but countably many $\delta > 0$, the δ-metric-thinning map $\theta^\delta : \mathbb{M} \to \mathbb{M}$ is continuous μ almost everywhere.
In particular, if Π_n weakly converges to Π, then $\theta^\delta(\Pi_n)$ weakly converges to $\theta^\delta(\Pi)$.

To prove the lemma, simply note that any δ such that $B_G(0, \delta)$ is a stochastic continuity set *for* Π_0 works.

Proof of Theorem 30. We prove (1), and then show how to reduce (2) to (1).
By increasing S if necessary, we can assume $\operatorname{diam}(S) < R$.
Denote the distributions of Π_n and Π by μ_n and μ respectively.
We call a factor graph $\mathscr{G}$ a *μ-continuity factor graph* if it has the property that

$$\lim_{n\to\infty} \overrightarrow{\mu_0}^n(\mathscr{G}) = \overrightarrow{\mu_0}(\mathscr{G}).$$

The same tedious technique used to prove Proposition 16 shows that factor graphs of the form $\mathscr{G}_{A,V} = (A \times V) \cap \overrightarrow{\mathbb{M}_0}$, where $A \subseteq \mathbb{M}_0$ is a μ_0 continuity set and $V \subseteq G$ is a bounded stochastic μ continuity set, are μ-continuity factor graphs.

The idea of the proof is that we will take a cheap graphing $\mathscr{G}$ for the limit process μ, and use it to produce a cheap μ_0-continuity graphing $\mathscr{H}$. The continuity property then gives us information about the costs of μ_n, *but only if we can ensure $\mathscr{H}$ is connected on* Π_n. The coarse density is exactly what allows us to ensure this.

Note that by outer regularity of the measure $\overrightarrow{\mu_0}$, for every factor graph $\mathscr{G}$ and $\varepsilon > 0$ there exists an *open* factor graph $\mathscr{G}' \supseteq \mathscr{G}$ such that $\overrightarrow{\mu_0}(\mathscr{G}') \leq \overrightarrow{\mu_0}(\mathscr{G}) + \varepsilon$. Therefore in the definition of cost one can replace "measurable graphing" by "open graphing".

Claim. Every *open* graphing $\mathscr{G}$ of μ contains a μ-continuity factor graph $\mathscr{H}_N$ such that $\mathscr{H}_N^N \supseteq \overrightarrow{\mathbb{M}_0} \cap (H_\delta \times S)$.

Here H_δ denotes the space of δ separated configurations as in Lemma 6.

Note that this condition and the assumption on R and S implies that $\mathscr{H}$ is *connected* on any δ-uniformly separated and R coarsely dense input. In particular, $\mathscr{H}(\Pi_n)$ is connected for every n.

To finish the proof from here: for any $\varepsilon > 0$, choose a graphing $\mathscr{G}$ of Π such that $\overrightarrow{\mu_0}(\mathscr{G}) \leq \text{cost}(\Pi) + \varepsilon$. Take $\mathscr{H}$ as in the above procedure. Then:

$$\begin{aligned}
\limsup_{n\to\infty} \text{cost}(\Pi_n) &\leq \limsup_{n\to\infty} \overrightarrow{\mu_0}^n(\mathscr{H}) && \text{As } \mathscr{H} \text{ is a graphing of } \Pi_n \\
&= \overrightarrow{\mu_0}(\mathscr{H}) && \text{Since } \mathscr{H} \text{ is a } \overrightarrow{\mu_0}\text{-continuity graphing} \\
&\leq \overrightarrow{\mu_0}(\mathscr{G}) && \text{As } \mathscr{H} \subseteq \mathscr{G} \\
&\leq \text{cost}(\Pi) + \varepsilon && \text{By assumption on } \mathscr{G}.
\end{aligned}$$

Since $\varepsilon > 0$ was arbitrary, this proves the result.

We must now prove the claim.

Recall from Lemma 8 that $\mathbb{M}_0$ and G admit *topological bases* $\{A_i\}$ and $\{V_j\}$ consisting of μ_0-continuity sets and μ stochastic continuity sets (respectively). So by definition of the subspace topology, $\overrightarrow{\mathbb{M}_0}$ admits a topological basis $\{\mathscr{G}_{A_i,V_j}\}$ consisting of μ-continuity factor graphs.

For each $k \in \mathbb{N}$ define

$$\mathscr{H}_k = \bigcup_{\substack{i,j\leq k \\ \mathscr{G}_{A_i,V_j}\subseteq\mathscr{H}}} \mathscr{G}_{A_i,V_j}.$$

Since $\mathscr{H}_k$ consists of *finitely many* continuity factor graphs, it is itself a continuity factor graph. Each $\mathscr{H}_k$ is also open, and increases to $\mathscr{H}$ as k tends to infinity.

As $\mathscr{H}$ is generating, $\{\mathscr{H}_k^k\}_{k\in\mathbb{N}}$ forms an open cover of the *compact* space $\overrightarrow{\mathbb{M}_0}\cap(H_\delta\times S)$. In particular, there exists N such that $\mathscr{H}_N^N \supseteq \overrightarrow{\mathbb{M}_0} \cap H_\delta \times S$, proving the claim.

One sees that the essential feature in the above proof strategy was *compactness*, and therefore it remains true for Ξ-labelled point processes if Ξ is compact, as mentioned.

With this observation in hand, we can now deduce the (2) statement from the (1). We will tediously produce a weakly convergent sequence of point processes, where each term has the same cost as Π_n and the weak limit factors onto Π, thus has cost at most the cost of Π. This proves the statement.

Choose $\delta' < \delta$ as in Lemma 11 ss that δ' metric thinning satisfies

$$\theta^{\delta'}(\Pi_n) \text{ weakly converges to } \theta^{\delta'}(\Pi) = \Pi.$$

Now observe by label trickery (see 11) we can find $[0,1]$-labelled point processes Υ_n each isomorphic to the respective Π_n and such that their underlying pointset is $\theta^{\delta'}(\Pi_n)$.

Note that Υ_n *might not weakly converge*, but it will have subsequential weak limits. All such subsequential weak limits will be some kind of (possibly random) labelling of Π.

To see this, let $\pi : [0,1]^{\mathbb{M}} \to \mathbb{M}$ be the map that forgets labels. Thus $\pi(\Upsilon_n) = \theta^{\delta'}(\Pi_n)$. Since π is continuous, it preserves weak limits. Let Υ be any subsequential weak limit of Υ_n, along a subsequence n_k. Then

$$\pi(\Upsilon) = \lim_{k\to\infty} \pi(\Upsilon_{n_k}) = \lim_{k\to\infty} \theta^{\delta'}(\Pi_{n_k}) = \Pi.$$

Now let $\Theta_n(\Upsilon_n)$ be *the input/output versions* of the (δ', R)-Delone thickenings that exist from Proposition 10. Here we use that Υ_n are free actions. By input/output we mean you keep track of which points of the thickening are input and output, as in Definition 10. In particular,

$$\text{cost}(\Theta_n(\Upsilon_n)) = \text{cost}(\Upsilon_n) = \text{cost}(\Pi_n),$$

where the first equality holds because we took the input/output version of the thickening.

Let Υ' denote any subsequential weak limit of $\Theta_n(\Upsilon_n)$. Then Υ' factors onto Π, by a similar argument to the earlier one about forgetting certain labels. Putting this all together:

$$\limsup_{n\to\infty} \operatorname{cost}(\Pi_n) = \operatorname{cost}(\Theta_n(\Upsilon_n)) \leq \operatorname{cost}(\Upsilon') \leq \operatorname{cost}(\Pi),$$

where the final inequality holds because cost can only increase under factors. □

Remark 33. In the second part of the proof, one might want to replace label trickery by something like "each Π_n is isomorphic to a random Delone set Υ_n, which has subsequential weak limits, so choose one such limit Υ...", but then it's not clear what the cost of Υ has to do with the cost of Π. One would require the Delone-ification process to preserve weak limits in some sense in order to relate $\operatorname{cost}(\Upsilon)$ and $\operatorname{cost}(\Pi)$.

Remark 34. One could argue that Abért and Weiss themselves engage in label trickery in [AW13]: they always assume their action is continuous on a compact space.

Theorem 31. If Π is a free point process, then its cost is at most the cost of the Poisson point process on G.

Proof. It is enough to prove that the IID Poisson point process $[0,1]^{\mathscr{P}}$ maximises the cost. This implies $\operatorname{cost}(\mathscr{P}) \leq \operatorname{cost}([0,1]^{\mathscr{P}})$, and $\operatorname{cost}([0,1]^{\mathscr{P}}) \leq \operatorname{cost}(\mathscr{P})$ by factoring, so $\operatorname{cost}([0,1]^{\mathscr{P}}) = \operatorname{cost}(\mathscr{P})$.

We know that Π weakly factors onto $[0,1]^{\Pi}$, and $[0,1]^{\Pi}$ factors onto the IID Poisson. We would like to say "so Π weakly factors onto the IID Poisson, and hence has less cost by the cost monotonicity statement", but we do not have either of these statements.

C'est la vie.

Note that Π is *abstractly isomorphic* to a δ uniformly separated process Π' Proposition 21. Then Π' is also free and has the same cost as Π. Now Π' weakly factors onto its own IID, more explicitly, there is a sequence of factors *labellings* $\Phi_n(\Pi')$ weakly converging to $[0,1]^{\Pi'}$. Because $\Phi_n(\Pi')$ is a labelling of Π' it is itself free and uniformly separated. Putting it all together:

$$\begin{aligned}
\operatorname{cost}(\Pi) &= \operatorname{cost}(\Pi') && \text{As they are isomorphic actions}\\
&\leq \limsup_{n\to\infty} \operatorname{cost}(\Phi_n(\Pi')) && \text{Cost can only increase for factors}\\
&\leq \operatorname{cost}([0,1]^{\Pi'}) && \text{by Theorem 30}\\
&\leq \operatorname{cost}([0,1]^{\mathscr{P}}) && \text{As } [0,1]^{\Pi'} \text{ factors onto } [0,1]^{\mathscr{P}}.
\end{aligned}$$

□

Note that the Poisson point process and *any* $[0,1]$ IID labelled point process have *the same* cost. In practice, one would try to work with the IID Poisson point process (as it is most explicit) or the IID Poisson net (since it's a net).

3.3 Some fixed price one groups.

We wish the results of the previous section to prove that certain groups have fixed price one. The strategy that immediately comes to mind is that one would prove that the IID Poisson on G has fixed price one. Even on groups of the form $G \times \mathbb{Z}$, we are unable to do this directly.

Question 7. Can one *explicitly* construct for every $\varepsilon > 0$ factor graphs of the IID Poisson on $G \times \mathbb{Z}$ of edge measure less than $1 + \varepsilon$?

Instead, we use the weak factoring strategy to reduce the above problem to a much simpler one, where we *can* construct such factor graphs.

3.3.1 $G \times \mathbb{Z}$

Definition 35. Let Π be a point process on G. Its *vertical coupling* on $G \times \mathbb{Z}$ is $\Delta(\Pi) = \Pi \times \mathbb{Z}$.

Here $\Delta : \mathbb{M}(G) \to \mathbb{M}(G \times \mathbb{Z})$ is induced by the diagonal embedding of G into $G^{\mathbb{Z}}$. For this reason one might prefer to call $\Delta(\Pi)$ the *diagonal coupling*, but this terminology will not be suitable when we go to $G \times \mathbb{R}$.

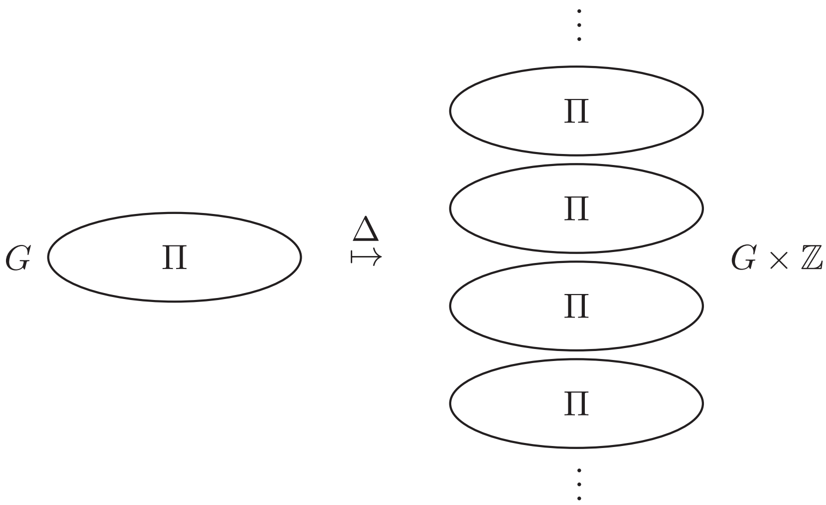

Figure 3.1: If you picture G as a pancake, then the point process Π is a scattering of blueberries on it. Its vertical coupling is a stack of pancakes with *the same* scattering of blueberries on every level. This generalises to other berries.

Lemma 12. The IID version $[0,1]^{\Delta(\Pi)}$ of a vertically coupled process has cost one.

The proof uses the fact that Bernoulli percolation of a factor graph can be implemented as a factor of IID. This sort of trick will be familiar to many, but we will nevertheless spell it out:

Definition 36. Let $\mathscr{G}$ be a factor graph of a point process Π. Its ε *edge percolation* is the factor graph $\mathscr{G}_\varepsilon$ defined on $[0,1]^\Pi$ in the following way: for points $g, h \in [0,1]^\Pi$ let

$$g \sim_{\mathscr{G}_\varepsilon} h \text{ whenever } g \sim_{\mathscr{G}} h \text{ and } \xi_g \oplus \xi_h < \varepsilon.$$

Here $\oplus$ denotes addition of the labels modulo one.

Observe that if (g, h_1) and (g, h_2) are edges of $\mathscr{G}(\Pi)$, then the random variables $\xi_g \oplus \xi_{h_1}$ and $\xi_g \oplus \xi_{h_2}$ are independent uniform once again.

Remark 35. If $\mathscr{G}$ is already a factor graph defined on $[0,1]^\Pi$, then we can implement $\mathscr{G}_\varepsilon$ on $[0,1]^\Pi$, that is, without adding further randomness (via the replication trick).

Proof of lemma. Let $\mathscr{G}$ be any graphing of Π with finite edge density. We *lift* it to a factor graph $\mathscr{G}^\Delta$ of $\Delta(\Pi)$ in the following way:

$$(g,n) \sim_{\mathscr{G}^\Delta(\Pi)} (h,n) \text{ if and only if } g \sim_{\mathscr{G}(\Pi)} h,$$

that is, as $\Delta(\Pi)$ is just copies of Π stacked on every level of $G \times \mathbb{Z}$, then we simply copy $\mathscr{G}$ onto every level of $G \times \mathbb{Z}$ as well.

Let $\mathscr{V}$ denote the factor graph of $\Delta(\Pi)$ consisting of *vertical* edges, that is for every $(g,n) \in \Delta(\Pi)$ we have an edge to $(g, n+1)$.

One can see that $\mathscr{V} \cup \mathscr{G}^\Delta$ is a connected factor graph. But this is also true when we percolate the edges level wise, that is, when we consider $\mathscr{V} \cup \mathscr{G}^\Delta_\varepsilon$. This is because if $(g,n) \sim_{\mathscr{G}^\Delta} (h,n)$ is an edge destroyed in the percolation, then we can slide up along vertical edges and consider the edge $(g,n+1) \sim_{\mathscr{G}^\Delta} (h,n+1)$ instead. Its chance of survival in the percolation is independent from the previous edge, and hence we get another go to cross over. By sliding up far enough we are guaranteed to be able to cross.

□

Remark 36. Recall from Section 1.3.5 that in order to prove $[0,1]^\Pi$ weakly factors onto $[0,1]^\Upsilon$, it suffices to prove that $[0,1]^\Pi$ weakly factors onto Υ.

Proposition 18. The IID Poisson on $G \times \mathbb{Z}$ weakly factors onto the vertically coupled Poisson of G.

Proof. We will construct factor maps $\Phi^n : [0,1]^\mathbb{M} \to \mathbb{M}$ that "straighten" the input in the following way: for a given input $\omega \in [0,1]^\mathbb{M}$, we select a "sparse" subset of its points. At each one of these we *propagate* them upwards by placing copies of them on the levels above. This will converge to a vertically coupled process for suitable inputs.

More precisely, let Π denote the (unit intensity) IID Poisson on $G \times \mathbb{Z}$. We will denote points of Π by $(g,l) \in G \times \mathbb{Z}$, and write $\Pi_{g,l}$ for its label.

We now define the factor map Φ^n in two stages as a thinning and then a thickening to simplify the analysis. Let

$$\Pi^{1/n} = \{(g,l) \in \Pi \mid \Pi_{g,l} \leq \frac{1}{n}\},$$

and $F_n = \{0\} \times \{0, 1, \ldots, n-1\}$. Set

$$\Phi^n(\Pi) = \Theta^{F_n}(\Pi^{1/n}).$$

Let us explain what this means:

- At the first step $\Pi \mapsto \Pi^{1/n}$, we independently thin Π to get a subprocess of intensity $\frac{1}{n}$. By the discussion in Example 6, the resulting process $\Pi^{1/n}$ is simply a Poisson point process on $G \times \mathbb{Z}$ of intensity $\frac{1}{n}$. We refer to the points of $\Pi^{1/n}$ as *progenitors*, as they're about to get busy.

- Each progenitor (g,l) places additional points with the same G-coordinate on the next $n-1$ levels above it. This is the map $\Pi^{1/n} \mapsto \Theta^{F_n}(\Pi^{1/n}) = \Phi^n(\Pi)$.

- By the discussion at Example 7, $\Phi^n(\Pi)$ is a process of unit intensity.

We will employ the following strategy to show that $\Phi^n\Pi$ weakly converges to the vertical Poisson:

1. The sequence $(\Phi^n\Pi)$ admits weak subsequential limits, which a priori might be random counting measures,
2. these subsequential limits are actually simple point processes,
3. all of these subsequential limits are vertical processes, and
4. that process is the vertical Poisson.

Recall that if (x_n) is a relatively compact sequence and every subsequential limit of (x_n) is x, then x_n converges to x.

By this basic fact and the above items, we can conclude that $\Phi^n\Pi$ weakly converges to the vertical Poisson.

We now verify that $\{\Phi^n(\Pi)\}$ is *uniformly tight*, proving (1). It suffices to verify that the distributions of point counts $N_C(\Phi^n\Pi)$, where $C = B_G(0,r) \times [L]$ denotes a cylinder whose base is a ball of radius r and its height (in levels) is L, are uniformly tight.

Let X_i denote the number of points in $B_G(0,r) \times \{i\}$ with label $\Pi_{g,i} \leq \frac{1}{n}$, that is, the number of progenitors on the ith level. Thus the X_i are IID Poisson random variables with parameter $\frac{\lambda(B_G(0,r))}{n}$.

One can explicitly describe the random variable $N_C(\Phi^n\Pi)$ in terms of the X_is, but for our purposes it is enough to observe that:

$$N_C(\Phi^n\Pi) \leq L \sum_{i=1}^{n} X_i.$$

The sum of independent Poisson random variables is again Poisson distributed (with parameter the sum of the parameters of the individual Poissons), so we see that $N_C(\Phi^n\Pi)$ is bounded in terms of a random variable *that does not depend on n*. Therefore $\{\Phi^n\Pi\}$ is uniformly tight.

To prove item (2) on the docket, note that the above shows that the point counts in $B_G(0,r) \times \{0\}$ for $\Phi^n\Pi$ are *exactly* Poisson distributed with parameter $\lambda(B_G(0,r))$. Thus if Υ is any subsequential weak limit of $\Phi^n\Pi$ and r is such that $B_G(0,r) \times \{0\}$ is a stochastic continuity set for Υ, then $N_{B_G(0,r)\times\{0\}}\Upsilon$ will also be Poisson distributed. In particular, Υ must be a *simple* point process.

For item (3), let Υ be any subsequential weak limit of $\Phi^n\Pi$. Observe that Υ is *vertical* if and only if $(g,l) \in \Upsilon$ implies $(g, l+1) \in \Upsilon$. The idea is that this property is satisfied for most points of $\Phi^n\Pi$, and therefore must be preserved in the weak limit. Note that a process is vertical if and only if its Palm measure is vertical almost surely.

We can now explicitly describe the Palm measure of $\Phi^n(\Pi)$. Recall from Theorem 13 that the Palm version $\Pi_0^{1/n}$ of $\Pi^{1/n}$ is simply $\Pi^{1/n} \cup \{(0,0)\}$.

To express the Palm version of the F_n-thickening of $\Pi^{1/n}$ (a la Example 16), it will be useful to introduce the following notation. For each $k \in NN$, let

$$\Pi_k^{1/n} = \Pi_0^{1/n} \cdot (0,1) = \{(g, l+1) \in G \times \mathbb{Z} \mid (g,l) \in \Pi_0^{1/n}\}.$$

That is, you simply shift $\Pi_0^{1/n}$ up by one level. Then $\Phi^n(\Pi_0) = \Pi_0^{1/n} \cup \Pi_1^{1/n} \cdots \cup \Pi_{n-1}^{1/n}$.

Denote by K a random integer chosen uniformly from $\{0, 1, \dots, n-1\}$. Then the Palm version of $\Phi^n(\Pi)$ is

$$\Phi^n(\Pi)_0 = \Pi^{1/n}_{-K} \cup \Pi^{1/n}_{-K+1} \cup \cdots \cup \Pi^{1/n}_{-K+n-1}$$

Let us say that a rooted configuration $\omega \in \mathbb{M}_0(G \times \mathbb{Z})$ *has an ε-successor* if there is a point approximately above the root $(0,0)$ in ω. More precisely, we define an *event*

$$\{\omega \text{ has an } \varepsilon\text{-successor}\} := \{\omega \in \mathbb{M}_0 \mid N_{B_G(0,\varepsilon)\times\{1\}}\omega > 1\}.$$

From this, we see

$$\mathbb{P}[\Phi^n(\Pi)_0 \text{ has an } \varepsilon\text{-successor}] \geq \frac{n-1}{n},$$

as $\Phi^n(\Pi)_0$ certainly has an ε-successor whenever $K < n-1$.

We've been assuming Υ is a subsequential weak limit of $\Phi^n(\Pi)$, and now we must unfortunately make this explicit: fix a subsequence n_i such that $\Phi^{n_i}(\Pi)$ weakly converges to Υ.

Choose a sequence ε_k tending to zero such that $B_G(0, \varepsilon_k) \times \{1\}$ is a stochastic continuity set for Υ. This is possible by Lemma 8. Then for each k

$$\frac{n_i - 1}{n_i} \leq \mathbb{P}[\Phi^{n_i}(\Pi)_0 \text{ has an } \varepsilon_k\text{-successor}] \to \mathbb{P}[\Upsilon_0 \text{ has an } \varepsilon_k\text{-successor}],$$

So Υ_0 has ε_k-successors almost surely for every k, and hence has 0-successors. That is, Υ is a vertical process, at last proving item (3).

Finally, for item (4) we observe that any vertical process is completely determined by its intersection with $G \times \{0\}$. We observed in the proof of item (2) that Υ is a Poisson point process on the 0th level, so it must be the vertical Poisson, as desired. $\square$

Corollary 3. Groups of the form $G \times \mathbb{Z}$ have fixed price one.

Proof. By the previous proposition and Remark 36, we know that the IID Poisson weakly factors onto the IID of the vertically coupled Poisson. Explicitly, there exists factor maps $\Psi^n : [0,1]^{\mathbb{M}} \to [0,1]^{\mathbb{M}}$ such that

$$\Psi^n\Pi \text{ weakly converges to } [0,1]^{\Delta(\mathscr{P})},$$

where Π is the IID Poisson on G and[1] $\mathscr{P}$ is the Poisson on G.

Choose $\delta < 1$ as in Lemma 11 such that metric δ-thinning preserves the weak limit. Note that because $\delta < 1$, the thinning commutes with the vertical coupling: that is, $\theta^\delta(\Delta(\mathscr{P})) = \Delta(\theta^\delta \mathscr{P})$. Therefore

$$\theta^\delta(\Psi^n\Pi) \text{ weakly converges to } [0,1]^{\Delta(\theta^\delta(\mathscr{P}))}.$$

Putting this all together,

$$\begin{aligned} \operatorname{cost}(\Pi) &\leq \limsup_{n\to\infty} \operatorname{cost}(\theta^\delta(\Psi^n\Pi)) && \text{As cost can only increase under factors} \\ &= \operatorname{cost}([0,1]^{\Delta(\theta^\delta(\mathscr{P}))}) && \text{by Theorem 30} \\ &= 1 && \text{by Lemma 12.} \end{aligned}$$

Since the IID Poisson has maximal cost, this proves that $G \times \mathbb{Z}$ has fixed price one. $\square$

[1] This is a slight abuse of notation: we were using $\mathscr{P}$ to denote the *law* of the Poisson point process, but in the above expression we treat it as if it were a random variable. We do this to prevent the profusion of asterisks representing pushforwards of measures.

Remark 37. By making further percolation-theoretic assumptions on G, one can *directly* show that $\text{cost}(\Phi^n(\Pi)) \leq 1 + \varepsilon_n$, where ε_n tends to zero. This is by constructing factor graphs on $\Phi^n(\Pi)$.

By using the Poisson net, one can prove an analogue of the Babson and Benjamini theorem [BB99] and show that the distance $\mathscr{D}_R$ factor graph on the Poisson point process on a *compactly presented* and one-ended group has a *unique* infinite connected component if R is sufficiently large.

Now on $\Phi^n(\Pi)$, we construct a factor graph as follows: add in all vertical edges, and the $\mathscr{D}_R$ edges horizontally. Now percolate the horizontal edges. One can show that by adding a small amount of edges to this, the result is a graph with a unique infinite connected component.

3.3.2 $G \times \mathbb{R}$

We now outline the modifications required to extend the $G \times \mathbb{Z}$ case to the following theorem:

Theorem 32. Groups of the form $G \times \mathbb{R}$ have fixed price one.

In the previous section, we asked of you to visualise $G \times \mathbb{Z}$ as an infinite stack of pancakes. Now we ask you to visualise $G \times \mathbb{R}$ as a forbidding infinite sausage (American readers may prefer to visualise $G \times \mathbb{R}$ as a pancake-dipped sausage).

Proof. The strategy will be exactly the same as in Proposition 18.

We define factor maps Φ^n of the IID Poisson Π using the same formula as in the $G \times \mathbb{Z}$ case. We claim these weakly converge to a point process Υ which is *vertical* in the sense that $(g, t) \in \Upsilon$ implies $(g, t+n) \in \Upsilon$ for all $n \in \mathbb{Z}$.

First we show $\{\Phi^n(\Pi)\}$ is uniformly tight. This works exactly as in the $G \times \mathbb{Z}$ case, except instead of counting progenitors X_i on $G \times \{i\}$, we count them on $G \times [i, i+1)$ for $i \in \mathbb{Z}$.

Next we show that any subsequential weak limit Υ of $\{\Phi^n(\Pi)\}$ is not just a random counting measure, but an actual point process. This follows as in the $G \times \mathbb{Z}$ case, as $\Phi^n(\Pi)$ has the same distribution in $G \times [0, 1)$ as a Poisson point process on $G \times \mathbb{R}$.

The proof that Υ is a vertical point process works the same as in the $G \times \mathbb{Z}$ case.

At this point one can observe that a vertical process is determine by its intersection with $G \times [0, 1)$, and therefore $\Phi^n(\Pi)$ weakly converges to a unique point process Υ.

We now adapt Lemma 12 to this context, and show that if Υ is *any* vertical point process, then its IID $[0, 1]^\Upsilon$ has cost one.

Let $\pi : G \times \mathbb{R} \to G$ denote the projection map. Observe that *if* Υ *is vertical*, then $\pi(\Upsilon)$ is discrete, and hence defines a point process on G. For contrast, observe that the projection $\pi(\Pi)$ of the Poisson point process Π is almost surely dense, and hence does not define a point process on G.

Let[2] $\mathscr{G}$ be a finite cost graphing of $\pi(\Upsilon)$. We lift this to a factor graph of Υ in the following way:

$$(g_1, t_1) \sim_{\mathscr{H}(\Upsilon)} (g_2, t_2) \text{ when } g_1 \sim_{\mathscr{G}(\pi(\Upsilon))} g_2 \text{ and } |t_1 - t_2| < 1.$$

[2]technically we are assuming such a thing exists, which will be the case if $\pi(\Upsilon)$ is a free point process on G, or a lattice shift where the lattice is finitely generated

Let $\mathscr{V}(\Upsilon)$ denote the set of *vertical edges*, that is

$$\mathscr{V}(\Upsilon) = \{((g,t),(g,t+n)) \in \Upsilon \times \Upsilon \mid n \in \mathbb{Z}\}.$$

Then as in Lemma 12, the vertical edges $\mathscr{V}(\Upsilon)$ together with an ε-percolation of $\mathscr{H}(\Upsilon)$ defines a cheap connected factor graph of Υ.

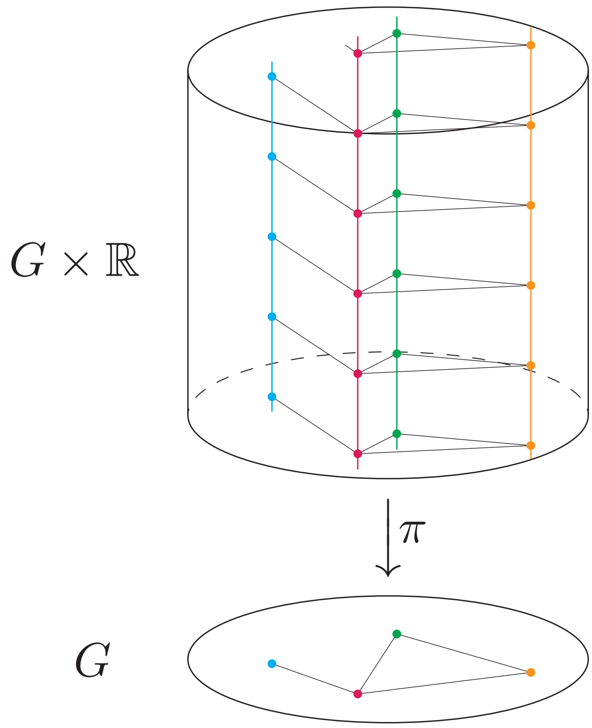

Figure 3.2: A portion of a graphing on the projection of a vertical process, and how it might look when lifted. Note that it gets wobbled a bit in the process.

We conclude from this that $G \times \mathbb{R}$ has fixed price one by the same kind of reasoning as in Corollary 3. □

Remark 38. The limiting process here is a Γ-periodic Poisson (see Example 4), where $\Gamma = \{0\} \times \mathbb{Z} < G \times \mathbb{R}$.

3.3.3 Groups containing noncompact amenable normal subgroups

Question 8. Suppose G admits a noncompact amenable normal subgroup N. Does G have fixed price one?

This is what one would expect from the discrete case.

I offer the skeleton of a proof, modelled on the $G \times \mathbb{Z}$ case. Of course, N takes the role of $\mathbb{Z}$. I will write as if N is nondiscrete, if it is discrete than just use $\{0,1\}^N$ instead of $\mathbb{M}(N)$, and Bernoulli percolation instead of the Poisson point process, and so on.

Proposition 19. Suppose G admits a noncompact amenable normal subgroup N. Then there exists a point process Υ on G such that $[0,1]^\Upsilon$ has cost one.

Proof. Fix Haar measures $\lambda_G, \lambda_N, \lambda_{G/N}$ on $G, N, G/N$ respectively.

For $a \in G$, denote by λ_{aN} the measure on $aN \subset G$ given by $\lambda_{aN}(V) = \lambda_N(a^{-1}V)$.

Let Π be the Poisson point process on G/N. We define a *random measure* $\eta(\Pi)$ on G by

$$\eta(\Pi) = \sum_{aN \in \Pi} \lambda_{aN}.$$

Let Υ be the Poisson point process[3] sampled from $\eta(\Pi)$.

To spell this out a bit more: we sample from a Poisson point process Π on G/N. This can be viewed as an ensemble of closed subsets in G, each of which are translates of N. On each of these we independently sample from a Poisson point process on N.

Observe that since Υ is a Poisson on every coset aN, it follows that if Υ_0 is a sample from the Palm measure of Υ, then $\Upsilon_0\big|_N$ is the Palm measure of the Poisson on N.

In particular, it defines a *hyperfinite* equivalence relation.

Let $\mathscr{V}$ denote a graphing of the Palm measure of the Poisson point process on N which is a directed $\mathbb{Z}$. This is possible from our discussion of the interpretation of amenability for point processes. This $\mathscr{V}$ will play the role of the vertical edges from the $G \times \mathbb{Z}$ case.

Now fix a graphing $\mathscr{G}$ of the Poisson point process on G/N. We lift this to a factor graph $\mathscr{H}$ of Υ in a similar way to what we did for the $G \times \mathbb{R}$ case. Let $\pi : G \to G/N$ denote the natural projection $g \mapsto gN$, and observe that $\pi(\Upsilon) = \Pi$. Define edges between points $g, h \in \Upsilon$ by

$$g \sim_{\mathscr{H}(\Upsilon)} h \text{ when } gN \sim_{\mathscr{G}(\pi(\Upsilon))} hN \text{ and } d(g,h) \leq d(g, hn) \text{ for all } n \in N \text{ such that } hn \in \Upsilon.$$

We combine these factor graphs as possible: on each coset $aN \in \Upsilon$, place down a copy of $\mathscr{V}(aN)$. Take the union of this with $\mathscr{H}(\Upsilon)$. This is a connected graph: every pair of points $g, h \in \Upsilon$ has a path $a_1N, a_2N, \ldots, a_kN$ in $\mathscr{G}(\pi(\Upsilon))$ between their projections gN and hN, the connected

Finally, $[0,1]^\Upsilon$ has cost one: take the above graphing and percolate the $\mathscr{H}$ edges. ☐

Let Π denote the IID Poisson on G. We now define a sequence of point process factors $\Phi^n(\Pi)$ which seem like they ought to weakly converge to Υ above.

Fix a sequence $F_n \subseteq N$ of Følner sets, and let $v_n = \lambda_N(F_n)$ denote their volumes. Let $\Psi^n : [0,1] \to \mathbb{M}(F_n)$ be a map such that $\Psi^n_*(\texttt{Leb})$ is the unit intensity Poisson point process on F_n.

If $g \in \Pi$, then we write Π_g for its uniform $[0,1]$ label.

We now define Φ^n as follows:

$$\Phi^n\Pi = \bigcup_{g \in \Pi | \Pi_g < v_n} g\Psi^n(\Pi_g).$$

That is, a sparse set of progenitors spreads out around itself N-wise a random sample of the Poisson from the Følner sets.

[3]This is an example of a *Cox process*: that is, a Poisson point process sampled according to a random measure

3.3.4 The next step

Definition 37. Let H be an lcsc group, and $t \in H$. We say t *generates a discrete* $\mathbb{Z}$ if t has infinite order and $\langle t \rangle \in \mathbb{M}(H)$.

The ultimate goal is to prove a statement approximately of the following form:

Conjecture. Suppose $G \times H$ is a product of noncompact but compactly generated groups, where H contains an element t generating a discrete $\mathbb{Z}$. Then $G \times H$ has fixed price one.

A particularly interesting case would be $\mathrm{SL}(2, \mathbb{R} \times \mathrm{SL}(2, \mathbb{R})$, where such a theorem would imply new vanishing results for the rank gradient of Farber sequences of lattices.

The idea was to try to apply the propagation technique from the $G \times \mathbb{Z}$ case. This weak limits on a periodic Poisson, but with respect to $\{0\} \times \langle t \rangle < G \times H$. I am unable to show that this process has cost one (the lifting technique does not work, as this subgroup has infinite covolume).

3.4 Rank gradient of Farber sequences vs. cost

Definition 38. Let (Γ_n) denote a sequence of lattices in a fixed group G.

The sequence is *Farber* if for every compact neighbourhood of the identity $V \subseteq G$ we have

$$\mathbb{P}[a\Gamma_n a^{-1} \cap V = \{e\}] \to 1 \text{ as } n \to \infty,$$

where $a\Gamma_n$ denotes a coset of Γ_n chosen randomly according to the (normalised) finite G-invariant measure on G/Γ_n.

Note that $a\Gamma_n a^{-1}$ is exactly the stabiliser of $a\Gamma_n$ for the action $G \curvearrowright G/\Gamma$. Thus the Farber condition says that the action on most points of the quotient is locally injective.

Equivalently, the condition states that $a\Gamma_n \cap Va = \{a\}$ with high probability. It is this second form that we will actually use in the proof below. We think of a as being a point sampled randomly from a fundamental domain for Γ_n in G, and thus it states that the V-neighbourhood around this point a meets the lattice shift $a\Gamma_n$ only trivially.

Definition 39. Let (Γ_n) denote a sequence of lattices in a fixed group G. Its *rank gradient* is

$$\mathrm{RG}(G, (\Gamma_n)) = \lim_n \frac{d(\Gamma_n) - 1}{\mathrm{covol}\, \Gamma_n},$$

whenever this limit exists.

Remark 39. If G is discrete, then the Γ_n are all finite index subgroups. The Nielsen-Schreier formula

$$\frac{d(\Gamma_n) - 1}{[G : \Gamma_n]} \leq d(G) - 1$$

shows that the terms in the rank gradient are at least bounded.

Gelander proved [Gel11] an analogue of this formula for lattices in connected semisimple Lie groups without compact factors.

In the Seven Samurai paper [ABB+17], it is shown that if G is a centre-free semisimple Lie group of higher rank with property (T), then *any* sequence of irreducible lattices (Γ_n) in G is automatically Farber, as long as $\mathrm{covol}(\Gamma_n)$ tends to infinity.

In the particular case of a decreasing *chain* $\Gamma = \Gamma_1 > \Gamma_2 > \ldots$ of finite index subgroups, Abért and Nikolov showed [AN12] that the rank gradient $RG(\Gamma, (\Gamma_n))$ can be descriped as the *groupoid* cost of an associated pmp action $\Gamma \curvearrowright \partial T(\Gamma, (\Gamma_n))$ on the boundary of a rooted tree.

The following discreteness assumption is needed for two reasons: one to ensure that certain processes will be *hard-core*, and thus to ensure that they have weak limits.

Definition 40. We say that a lattice $\Gamma < G$ is *δ-uniformly discrete* if all of its *right* cosets $\Gamma a \in \Gamma \backslash G$ are δ uniformly separated as subsets of G. That is, for all distinct pairs $\gamma_1, \gamma_2 \in \Gamma$, we have $d(\gamma_1 a, \gamma_2 a) \geq \delta$. Equivalently by left-invariance of the metric, $d(e, a^{-1}\gamma a) \geq \delta$ for all $\gamma \in \Gamma$ not the identity, and where $e \in G$ denotes the identity of G.

If (Γ_n) is a sequence of lattices, then we say it is δ uniformly discrete if each Γ_n is δ uniformly discrete in the above sense.

Theorem 33. Let (Γ_n) be a Farber sequence of *cocompact* lattices. Suppose further that the sequence is *uniformly discrete*. If its rank gradient exists, then

$$\mathrm{RG}(G, (\Gamma_n)) \leq \mathrm{cost}(G) - 1.$$

In particular, if G has fixed price one then the rank gradient vanishes.

The above theorem was proved independently by Carderi in [Car18], with a similar method but in a drastically different language (namely, that of ultraproducts of actions). The theorem is therefore his, but hopefully the following proof is more appealing to those who find ultra-things distasteful.

Proof of theorem. Recall that the groupoid cost of a lattice shift is

$$\mathrm{gcost}(G \curvearrowright G/\Gamma_n) = 1 + \frac{d(\Gamma_n) - 1}{\mathrm{covol}\,\Gamma_n},$$

which is essentially the term appearing in the rank gradient definition. We would therefore like to take a weak limit of these actions to get some free point process, and then appeal to the cost monotonicity result. Of course, this is completely senseless: the intensity of the lattice shift tends to zero, so it weak limits on the empty process.

Therefore we *thicken* the lattice shifts to get processes Π_n with a nontrivial weak limit. This thickening procedure must be done correctly, so that we can apply our (weak) cost monotonicity result.

We will produce a sequence of $[0,1]$-marked point processes Π_n such that

- each Π_n is a 2δ-net,
- each Π_n is a factor of the lattice shift $a\Gamma_n$, and so has cost *at least* $1 + \frac{d(\Gamma_n)-1}{\mathrm{covol}\,\Gamma_n}$, and
- they have a weak limit Υ with IID $[0,1]$ labels.

Then

$$\mathrm{RG}(G, (\Gamma_n)) + 1 \leq \lim_{n\to\infty} \mathrm{gcost}(\Pi_n) \leq \mathrm{cost}(\Upsilon) = \mathrm{cost}(G)$$

by the cost monotonicity result, as desired.

We view the space of *right* cosets $\Gamma_n \backslash G$ as compact metric spaces, where the distance between two cosets $\Gamma_n b_1, \Gamma_n b_2 \in \Gamma_n \backslash G$ is just their distance as closed subsets of G.

Let $B_n = \{\Gamma_n b_1^n, \Gamma_n b_2^n, \dots, \Gamma_n b_{k_n}^n\}$ be a collection of 2δ-nets in $\Gamma_n \backslash G$, where δ is the uniform discreteness parameter. We also choose $b_1 = e$.

This specifies a sequence of *thickenings* Θ_n of the corresponding lattice shifts: that is, $a\Gamma_n \mapsto aB_n$.

Note that $\Theta_n(a\Gamma_n)$ is a 2δ-net: it's true that $d(a\Gamma_n b_i^n, a\Gamma_n b_j^n) = d(\Gamma_n b_i^n, \Gamma_n b_j^n) \geq \delta$ for $i \neq j$ by our choice of B_n, and points of $\Gamma_n b_i^n$ are uniformly separated too exactly by our uniform discreteness assumption. It is also 2δ-coarsely dense, by the same property for B_n.

Since $\{\Theta_n(a\Gamma_n)\}$ is a collection of random 2δ-nets, it is automatically uniformly tight, and all subsequential weak limits are 2δ-nets (and in particular, simple point processes).

At this point we could consider the input/output version of Θ_n (so that the image process is weakly isomorphic to the original, and in particular has the same groupoid cost), and then pass to a subsequential weak limit. Our issue here is that one would have to demonstrate that this weak limit is a *free* action in order to compare its cost to the cost of the ambient group. To do this, one would have to use the Farber condition in an essential way.

We bypass this by a labelling trick: note that the IID of *any* point process is automatically a free action (as any two points of it will receive distinct values almost surely, killing any possible symmetries). So we will limit on an IID labelled process instead.

Let $(a\Gamma_n, \xi)$ denote the *periodic* IID lattice shift (see Remark 6). Note that this process has the same cost as the regular lattice shift. We will thicken as before, but this time distribute labels: let

$$\Theta_n(a\Gamma_n, \xi) = \bigcup_{i=1}^{k_n} a\Gamma_n b_i^n \times \{\xi_i\}.$$

Let Υ denote any subsequential weak limit of $\Theta_n(a\Gamma_n, \xi)$ and $\pi(\Upsilon)$ its unlabelled version, where $\pi : G \times [0,1] \to G$ is the projection map that forgets labels. Then the weak limit of $\pi(\Theta_n(a\Gamma_n, \xi)$ over the same set of indices as for Υ weakly converges to $\pi(\Upsilon)$, as the projection is a continuous map and hence preserves weak limits.

Our task is to show that $\Upsilon = [0,1]^{\pi(\Upsilon)}$.

The idea of the proof is the following: fix $C \subseteq G$ to be a bounded stochastic continuity set for Υ. We want to prove that the labels of the points of $\Theta_n(a\Gamma_n, \xi)$ in C are independent and uniform on $[0,1]$. They are already $\texttt{Unif}[0,1]$ by definition, so we must now consider their dependencies. Again, by definition, points of C arising from *distinct* $a\Gamma_n b_i^n$ are automatically independent. The only dependency issue that can arise is when $a\Gamma_n b_i^n \cap C$ has *multiple* points. We will show that this is a vanishingly rare event.

This will be achieved via the following lemma:

Lemma 13. Let $C \subseteq G$ be compact. If (Γ_n) is a Farber sequence and $B_n \subseteq G$ is any sequence of finite subsets, then $\mathbb{P}[\exists b \in B_n \text{ such that } |a\Gamma_n b \cap C| > 1] \to 0$.

Proof. Apply the Farber condition with any set V containing CC^{-1}. If $b \in B_n$ is such that there are $a\gamma_1, a\gamma_2$ *distinct* elements of $a\Gamma_n b \cap C$, then

$$(a\gamma_1 b)(a\gamma_2 b)^{-1} = a\gamma_1\gamma_2^{-1}a^{-1} \text{ is in } CC^{-1},$$

so $a\gamma_1\gamma_2^{-1}a^{-1}ab = a\gamma_1\gamma_2^{-1}b \in Vab$, and this element is also in $a\Gamma_n b$. By the Farber condition,

$$a\Gamma_n b \cap Vab = \{ab\}$$

with high probability, and so

$$\mathbb{P}[\exists b \in B_n \text{ such that } |a\Gamma_n b \cap C| > 1] \leq \mathbb{P}[a\Gamma_n a^{-1} \cap V = \{e\}] \to 0,$$

finishing the proof. □

Let $\boldsymbol{V} = (V_1, V_2, \ldots, V_k)$ denote a collection of bounded stochastic continuity sets for Υ, and $[0, \boldsymbol{p}) = ([0, p_1), [0, p_2), \ldots, [0, p_k))$ a family of intervals in $[0, 1]$. We denote by $\boldsymbol{V} \times [0, \boldsymbol{p}) = (V_1 \times [0, p_1), \ldots, V_k \times [0, p_k))$ the stochastic continuity set of $[0, 1]^\Upsilon$.

Let C be a compact set large enough to contain $\bigcup_i V_i$.

On the event from the lemma,

$$N_{\boldsymbol{V}}(\Theta_n(a\Gamma_n, \xi)) = N_{\boldsymbol{V}}([0, 1]^{\Theta_n(a\Gamma_n)}),$$

where by $\Theta_n(a\Gamma_n)$ we simply mean $(\Theta_n(a\Gamma_n, \xi))$ with the labels erased.

Therefore $\Theta_n(a\Gamma_n, \xi)$ converges weakly to $[0, 1]^\Upsilon$, finishing the proof.

□

Remark 40. The label trickery in the above proof seems odd, but step back a bit: suppose we could (somehow) establish that Υ is a free point process? What then? We would apply the point process version of Abért-Weiss and conclude that it... weakly factors onto its own IID. So why not aim for that IID right from the outset?

Remark 41. To extend the theorem to nonuniform lattices one would need to improve the cost monotonicity result. It's certainly true that if a sequence of point processes Π_n weakly converge to Π, then one can find a great many continuity factor graphs for the limit, but the issue is that when you *pull them back* to the processes Π_n, there is no guarantee that you get something connected (and thus telling you about the cost, that is, rank gradient).

3.5 Point processes on symmetric spaces

Much of the above extends to point processes on spaces of the form $X = G/K$, where $K \leq G$ is a compact subgroup. In particular, it applies when X is a Riemannian symmetric space and $G = \text{Isom}(X)$.

We wish to investigate how much of the point process theory we've developed on *groups* carries over to such homogeneous spaces G/K. After all, it is somewhat perverse to consider a point process on $\text{Isom}(\mathbb{H}^d)$ rather than on hyperbolic space $\mathbb{H}^d$ itself. For in the latter case one can – if sufficiently educated in these dark arts – actually have a picture in one's mind's eye. And you'll have a suite of geometric tools, to boot.

One fixes a G-invariant measure and a proper left-invariant metric on G/K (such a metric exists even without assuming G is Lie, see Proposition 2.4 of [AD13]). One defines point processes and Ξ-marked point processes as invariant probability measures for $G \curvearrowright \mathbb{M}(X)$ and $G \curvearrowright \Xi^{\mathbb{M}(X)}$ respectively. Intensity is defined just as before, and is well-defined for the same reason. And our old friend the Poisson point process and his IID cousin can be defined just the same on X as on G.

Let us introduce the following space

$$\mathbb{M}_K(X) = \{\omega \subseteq X \mid K \in \omega\}$$

of configurations in X rooted at K.

Observe that a subset $A \subseteq \mathbb{M}_K(X)$ does *not* give a well-defined thinning $\theta^A : \mathbb{M}(X) \to \mathbb{M}(X)$. One would want to define for $\omega \in \mathbb{M}(X)$

$$\theta^A(\omega) = \{gK \in \omega \mid g^{-1}\omega \in A\},$$

but this is sensitive to your choice of representative. We see that the remedy is to take an additional quotient: note that $K \curvearrowright \mathbb{M}_K(X)$ *on the left* (the right action is trivial), so one can form a quotient $K \setminus \mathbb{M}_K(X)$. Ignoring measurability issues for now, one can verify that there is a correspondence between subsets $A \subseteq K \setminus \mathbb{M}_K(X)$ and equivariant thinnings $\theta : \mathbb{M}(X) \to \mathbb{M}(X)$.

We now address the measurability issue:

Lemma 14. The action $K \curvearrowright \mathbb{M}_K(X)$ is *smooth*. That is, the quotient space $K \setminus \mathbb{M}_K(X)$ is a standard Borel space.

Proof. One can appeal to Corollary 2.1.13 of Zimmer's book [Zim84], which states that *any* continuous action by a compact group is smooth, however it is simple enough to prove directly in this case (the proof is very reminiscent of the section of Holroyd and Peres' paper [HP03] labelled "Extension to more general point processes").

We will construct a function $F : \mathbb{M}_0(X) \to [0, 1]$ with the property that $F(\omega) = F(\omega')$ if and only if $\omega' \in K\omega$. This is a characterisation of smoothness.

Fix a family U_n of open subsets of G with the property that it *separates* $\mathbb{M}_0(X)$ in the sense that $|\omega \cap U_n| = |\omega' \cap U_n|$ for all n if and only if $\omega = \omega'$.

Let $f : \{0, 1\}^{\mathbb{N}} \to [0, 1]$ be any continuous injection, and consider the map $F : \mathbb{M}_0(X) \to [0, 1]$ given by

$$F(\omega) = \inf_{k \in K} f((\mathbb{1}_{U_n \cap k \cdot \omega \neq \varnothing})_n)$$

Note that the component functions $\omega \mapsto \mathbb{1}_{U_n \cap \omega \neq \varnothing}$ are lower semicontinuous, so the infimum exists.

The function is constant on K-orbits by definition, but by the separating nature of the family $\{U_n\}$ it also takes distinct values for orbits in distinct orbits. □

This quotient $K \setminus \mathbb{M}_K(X)$ plays the same role for $\mathbb{M}(X)$ as $\mathbb{M}_0(G)$ played for $\mathbb{M}(G)$.

If Π is a point process on X with law μ, then we can define its Palm measure μ_0 on $K \setminus \mathbb{M}_K(X)$ as before: the measure of $A \subseteq K \setminus \mathbb{M}_K(X)$ is

$$\mu_0(A) = \frac{\text{intensity } \theta^A \Pi}{\text{intensity } \Pi}.$$

Note that the rerooting equivalence relation $\mathcal{R}$ on $K \setminus \mathbb{M}_K(X)$ given by

$$\omega \sim_{\mathcal{R}} g^{-1}\omega \text{ for all } gK \in \omega$$

where configurations should be understood mod K defines a *countable* Borel equivalence relation. By the same techniques as for G-processes one can show that this is pmp.

There does not appear to be a good groupoid structure on $K \setminus \mathbb{M}_K(X)$, so we will have to make do with the equivalence relation (and are probably best to only consider free point processes).

As before, factor graphs of Π correspond to Borel graphs on $(K \setminus \mathbb{M}_K(X), \mathcal{R})$.

Therefore we are able to define the X-*cost* of a point process Π on X.

At this point though, the immediate question should be: how does this X-cost relate to the cost of Π *as a G-action* $G \curvearrowright (\mathbb{M}(X), \mu)$? Recall that if $G \curvearrowright (Y, \mu)$ is a *free* pmp action of G on a standard Borel space, then we define its cost by choosing any isomorphism of $G \curvearrowright (Y, \nu)$ with an invariant point process on G, and taking the cost of that point process.

Suppose now that $K \leq G$ is Haar null. Then the Poisson point process Π_G on G equivariantly maps onto the Poisson point process Π_X on X via the projection map $G \mapsto X$ and the mapping theorem, see Section 2.3 of [Kin93].

Thus $\text{cost}_G \Pi_G \leq \text{cost}_G \Pi_X$, and so $\text{cost}_G \Pi_G = \text{cost}_G \Pi_X$ as the Poisson maximises cost.

Question 9. Is $\text{cost}_X(\Pi_X) = \text{cost}_G(\Pi_X)$?

This would give a more accessible way to prove fixed price.

3.6 Point processes vs. cross-sections

Let $G \curvearrowright (X, \mu)$ be a pmp action of an lcsc group on a standard Borel space X. Previous works (see [KPV15], [Car18], [CLM18] for example) have made use of the following concept:

Definition 41. A *lacunary cross-section* for the action $G \curvearrowright (X, \mu)$ is a Borel subset $Y \subseteq X$ with the property that $|Gx \cap Y|$ is *countable* for μ almost every $x \in X$.

- If $V \subseteq G$ is an identity neighbourhood with $Vx \cap Y = \{x\}$ for μ almost every $x \in X$, then we call Y a V-*lacunary cross-section*.
- If $C \subseteq G$ is a compact identity neighbourhood with $CY = X$ (μ almost sure), then we call Y a C-*cocompact cross-section*.
- If there is a non-empty open subset $U \subseteq G$ with $\mathbb{E}_\mu [\#\{g \in U \mid g^{-1}x \in U\}] < \infty$, then we say Y has *finite intensity*.

Theorem 34 (Forrest [For74], Proposition 2.10)**.** Suppose $G \curvearrowright (X, \mu)$ is an essentially free action on a standard probability space. Then the action admits a cocompact cross-section.

See Section 4 of [KPV15] for a modern proof of the above fact, and Section 3.B of [Kec] for more references.

Example 24. Suppose $G \curvearrowright (\mathbb{M}, \mu)$ is a point process. Then $\mathbb{M}_0$ is a lacunary cross-section (by definition). It has finite intensity if and only if the point process has finite intensity. It is V-lacunary if and only if the distance between distinct pairs of points in a random sample Π of μ avoid V. It is C-cocompact if and only if the C-neighbourhood of a random sample Π is all of G.

It appears to have gone unobserved that the above example is (essentially) *the only* example of a cross section. That is, suppose $G \curvearrowright (X, \mu)$ is a pmp action with $Y \subset X$ a choice of cross-section. Define

$$\mathscr{V} : (X, \mu) \to Y^{\mathbb{M}}$$
$$\mathscr{V}(x) = \{(g, y) \in G \times Y \mid x = gy\}.$$

This map is the orbit viewing map. You sample a μ-random point x, and look at its orbit Gx. This can be identified with G (as the action is free). The points of its orbit that meet Y form a countable subset of G by assumption. That is our resulting configuration. The purpose behind recording not just the g but also the particular $y \in Y$ with $x = gy$ is that the resulting map is *injective*. One can also check that the viewing map is G-equivariant. We therefore identify the action $G \curvearrowright (X, \mu)$ with $G \curvearrowright (y^{\mathbb{M}}, \mathscr{V}_* \mu)$. Moreover, the cross-section $Y \subset X$ is identified with ${\mathbb{M}_0}^Y \subset \mathbb{M}^Y$ under this map. Thus we have shown:

Proposition 20. Every free pmp action $G \curvearrowright (X, \mu)$ is isomorphic to a marked point process. This marked point process can be assumed to be coarsely dense and uniformly separated (that is, a random marked net). In particular, it can be assumed to have finite intensity.

If $\Gamma \curvearrowright (X, \mu)$ is a pmp action of a *discrete group* then X serves as a cross-section for the action, and the above simply describes the fact that we can view the action as an invariant X-colouring of Γ.

Here a distinction arises between discrete groups and nondiscrete groups: every free pmp action of a nondiscrete group is isomorphic to an *unlabelled* point process:

Proposition 21. Every free Ξ-marked point process is (abstractly) isomorphic to an unlabelled point process.

Proof. The method here is necessarily ad hoc, so we sketch the idea and let the reader fill in the blanks to their own level of satisfaction.

Suppose μ is a free Ξ-marked point process. We assume it is ergodic. Then using the method described at Lemma 3 we can find a factor map $\Phi : (\Xi^{\mathbb{M}}, \mu) \to \mathbb{M}$ such that $\Phi_* \mu$ is δ uniformly separated, for some $\delta > 0$. By a similar map to the orbit viewing map above, this lets us construct an isomorphism of μ with a uniformly separated Ξ'-marked point process, where $\Xi' = \Xi^{\mathbb{M}_0}$.

We wish to *locally encode* the labels of this process. To illustrate the idea, suppose we have a 10-separated point process Π on $\mathbb{R}^2$ which has labels from the set $\{+, -\}$. Construct a factor map of Π in the following way: if $(x, y) \in \Pi$ has label $-$, then delete the label and add the points $(x-1, y)$ and $(x+1, y)$. If its label is $+$, then delete the label add the points $(x-1, y), (x+1, y), (x, y-1), (x, y+1)$. In this way we have (fairly literally) encoded the label as a configuration of points – that is, the factor map to $\mathbb{M}$ is *injective*. See Figure 3.6.

In general, choose c much smaller than the uniform separation constant δ. Fix a Borel isomorphism I of the mark space Ξ' with the subspace of $\mathbb{M}_0(B_G(0, c))$ consisting of rooted configurations with 10 points (say). Finally, fix a nontrivial element g close to the identity.

Now define a factor map of $\Phi_* \mu$ in the following way: at each point x of the process delete its label ξ_x, and add in xg and $xgI(\xi_x)$. This map can be constructed in a way to be injective (one can identify which points were from the original process, and then reconstruct their label).

□

Remark 42. The preimage of a cross-section under an equivariant factor map is a cross section. Thus in essence, a choice of cross-section for a free pmp action is essentially the same thing as fixing a point process factor.

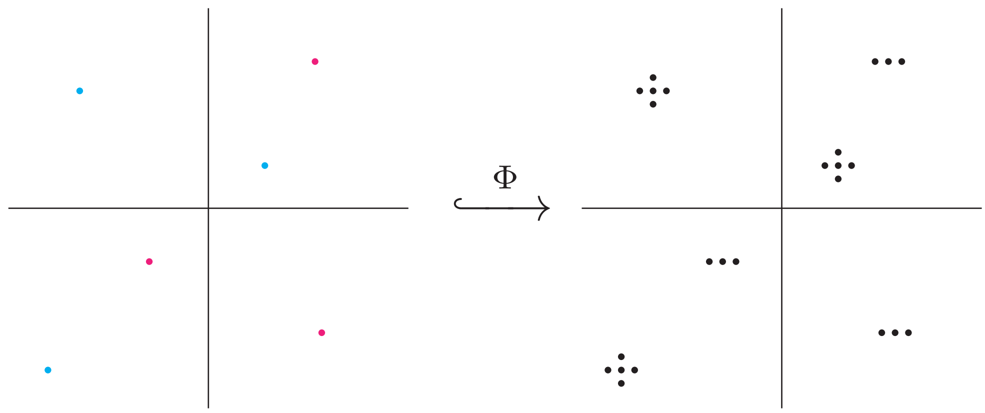

Figure 3.3: Here we view the cyan points as labelled by + and the magenta points as labelled by −. The map Φ spatially encodes the label near each point.

We have basically treated the existence of cross-sections as a black box. I have never *really* been able to understand what goes on inside it. Despite it not being logically necessitated, it is still nevertheless deeply upsetting that I am unable to come up with an alternative point process focused proof of the existence of cross-sections for free pmp actions $G \curvearrowright (X, \mu)$.

Our attempts have largely focused on the following idea: Given the action, we wish to produce an equivariant map $\Phi : (X, \mu) \to \mathbb{M}$. The idea is that one might be able to produce a map $\Psi : (X, \mu) \to \mathrm{Sub}(G)$, where $\mathrm{Sub}(G)$ denotes the space of closed subsets of G, and then produce a further equivariant selection map $\Sigma : \mathrm{Sub}(G) \to \mathbb{M}$. Neither of these steps are especially clear.

For the first, the following approach was suggested by Miklós Abért: fix a Borel isomorphism $I : X \to [0, 1]$. Then for each $x \in X$ define $\psi_x : G \to [0, 1]$ by

$$\psi_x(g) = I(g^{-1}x).$$

This is an equivariant way to associate measurable functions to points of X. One could then attempt to *smooth* ψ_x in some way, perhaps by taking pointwise averages over balls, to produce a *continuous* function $\psi'_x : G \to [0, 1]$. Then for each $t \in [0, 1]$ the level set

$$\{g \in G \mid \psi'_x(g) = t\}$$

defines a random closed subset of G. We are ignoring measurability issues for now.

If one is careful, then one can ensure that this random closed set is non-empty. An approach like this might reduce the first factoring problem to the second.

As for selecting points from this random closed set, I am at a loss. It seems difficult to do whilst maintaining equivariance.

In [KPV15], the following theorem is proved (see Proposition 4.3 of that paper) and described as folklore:

Theorem 35. Let G be a unimodular lcsc group, and $G \curvearrowright (X, \mu)$ an essentially free pmp action on a standard Borel space. Let $Y \subset X$ be a (cocompact) Borel cross-section. Fix a Haar measure λ for G. Then

1. The orbit equivalence relation of $G \curvearrowright (X, \mu)$ restricted to Y, that is,

$$\mathcal{R} = \{(y, y') \in Y \times Y \mid y \in Gy'\}$$

defines a cber.

2. The set $Z = \{(x, y) \in X \times Y \mid x \in Gy\}$ is Borel, and the projection $\pi_l : Z \to X$ onto the first coordinate is countable-to-one. Define a measure η on Borel subsets A of Z by
$$\eta(A) = \int_X \left|\pi_l^{-1}(x)\right| d\mu(x).$$
Then there exists a unique probability measure ν on Y and a number $0 < \operatorname{covol} Y < \infty$ such that
$$\Psi_*(\lambda \times \nu) = \operatorname{covol} Y \eta, \text{ where } \Psi : G \times Y \to Z \text{ is } \Psi(g, y) = (gy, y).$$

3. The probability measure ν is $\mathcal{R}$-invariant.

4. If $Y' \subset X$ is a different cross-section with corresponding equivalence relation $\mathcal{R}'$, then there exists Borel subsets $Y_0 \subseteq Y$ and $Y_0' \subseteq Y'$ and a Borel bijection $\alpha : Y_0 \to Y_0'$ such that
   - Y_0 meets almost every $\mathcal{R}$ orbit and Y_0' meets almost every $\mathcal{R}'$ orbit,
   - We have $\alpha_* \left(\nu\big|_{Y_0}\right) = \frac{\operatorname{covol} Y}{\operatorname{covol} Y'} \nu'\big|_{Y_0'}$.

5. $(Y, \mathcal{R}, \nu)$ is ergodic if and only if $G \curvearrowright (X, \mu)$ is ergodic,

6. $(Y, \mathcal{R}, \nu)$ is aperiodic if and only if G is noncompact,

7. $(Y, \mathcal{R}, \nu)$ is amenable if and only if G is amenable.

We give some comments on this theorem, and how it translates into statements about point processes.

1. This is mainly concerned with measurability issues which are clear in the point process case.

2. This is the analogue of the construction of the Palm measure and the edge measure.

3. This is verifying that the equivalence relation is pmp, as in Section 1.3.1.

4. This constructs an equivariant matching between the associated point process factors.

5. This is as in Section 1.3.2.

6. This is as in Section 1.1.2, where it is shown that point processes on noncompact groups consist of infinitely many points or zero points.

7. This is as in Section 1.5.

Given these comments, one may ask what the whole point of the enterprise was in the first place. The explanation is that I came across the above theorem and realised the connection *after* figuring out the point process versions when I was trying to understand the Holroyd and Peres paper [HP03]. For someone with my background, the proofs presented in this thesis are much more transparent, but that's not much solace given that the theorem is described as folklore.

3.6.1 Weak containment and the failure

Our initial motivation in studying point processes and cost was a paper of Abért and Weiss [AW13]. They proved that Bernoulli shift actions are *weakly contained* in all free pmp actions of discrete groups.

For a full description of weak containment, see the survey paper of Burton and Kechris [BK].

As far as the author is aware, weak containment of pmp actions for locally compact second countable (lcsc) groups has not been explored (at least publicly). The following is a tentative definition one could make:

Definition 42 (Tentative)**.** Let $G \curvearrowright^\alpha (X, \mu), G \curvearrowright^\beta (Y, \nu)$ be pmp actions of an lcsc group G. We say α *is weakly contained in* β (symbolically denoted by $\alpha \preceq \beta$) if for all Borel subsets $A_1, A_2, \ldots, A_n \subseteq X$, finite $F \subset G$, and $\varepsilon > 0$ there exists Borel subsets $B_1, B_2, \ldots, B_n \subseteq Y$ such that

$$|\mu(\gamma A_i \cap A_j) - \nu(kB_i \cap B_j)| < \varepsilon \text{ for all } i, j \in [n], \gamma \in F.$$

This is the usual definition when G is discrete. One could try replacing F by a compact set $K \subset G$, but this is no more powerful a definition: recall that the *weak topology* on $\operatorname{Aut}(X, \mu)$ is the topology generated by the maps $\{T \mapsto \mu(T(A) \triangle A) \mid A \subseteq X \text{ is Borel}\}$, and that a pmp action of G on (X, μ) is equivalent to a *continuous* homomorphism $G \to \operatorname{Aut}(X, \mu)$. By taking F sufficiently fine in K, one gets an equivalent definition.

One can also assume that the sets A_i form a partition of X, and B_i a partition of Y.

The viewpoint taken in [AW13] is that weak containment for discrete groups Γ is about symbolic dynamics representations of the actions. If $X = \bigsqcup_{i=1}^n A_i$ is a partition of X, then let $\Phi : X \to [n]^\Gamma$ denote the following map

$$\Phi_x(\gamma) = i, \text{ where } \gamma^{-1}x \in A_i.$$

This is a Γ-equivariant map, and so $\Phi_*\mu$ defines an invariant $[n]$-colouring on Γ.

Conversely, if $\Psi : X \to [n]^\Gamma$ is an equivariant map, then the sets $\{x \in X \mid \Psi_x(e) = i\}$ (where $e \in \Gamma$ denotes the identity) defines a measurable partition, and the associations are mutual inverses.

If Ξ is a complete and separable metric space, define

$$E(\alpha, \Xi) = \{\Phi_*\mu \in \operatorname{Prob}(\Xi^\Gamma) \mid \Phi : X \to \Xi^\Gamma \text{ is measurable and equivariant}\}$$

to be the space of all possible factor Ξ-colourings of a given action $\Gamma \curvearrowright^\alpha (X, \mu)$. Then the topology of weak convergence on $E(\alpha, \Xi)$ is itself a complete and separable metric space.

It can be shown that

Theorem 36 (Lemma 8 [AW13], Proposition 3.6 [TD15])**.** The following are equivalent for pmp actions $\Gamma \curvearrowright^\alpha (X, \mu)$ and $\Gamma \curvearrowright^\beta (Y, \nu)$ of a discrete group Γ:

- α is weakly contained in β,
- for all compact Ξ, $E(\alpha, \Xi) \subseteq \overline{E(\beta, \Xi)}$,
- for all $n \in \mathbb{N}$, $E(\alpha, [n]) \subseteq \overline{E(\beta, [n])}$,

- for $\kappa = \{0,1\}^{\mathbb{N}}$ (the Cantor space), $E(\alpha, \kappa) \subseteq \overline{E(\beta, \kappa)}$.

Here $\overline{E(\beta, \Xi)}$ denotes the topological closure of $E(\beta, \Xi)$ with respect to weak convergence.

The idea for extending the concept of weak containment was that a suitable replacement for Ξ^Γ when G is a nondiscrete group might be $\Xi^{\mathbb{M}}(G)$.

In this language what we prove as an analogue of the Abért-Weiss statement is that if Π is a free point process with law μ, then $[0,1]^\Pi \in \overline{E(\mu, [0,1])}$. To prove the full analogue of weak containemnet we would have to show that any *factor* of $[0,1]^\Pi$ can be achieved as a weak limit of factors of μ.

Question 10. Is weak factoring transitive? For example, if Π weakly factors onto Υ, and Υ factors onto Υ', then does Π weakly factor onto Υ'?

If μ_n weakly converges to μ, and $\theta : (\mathbb{M}, \mu) \to \mathbb{M}$ is a *thinning*, then one can find a subsequence n_k such that $\theta_* \mu_{n_k}$ weakly converges to $\theta_* \mu$. To do this, one notes that θ is determined by a subset $A \subseteq (\mathbb{M}_0, \mu)$. Then A can be approximated arbitrarily well by μ-continuity sets, and the corresponding factors will weakly converge.

Already we start to see issues. Why should $E(\alpha, \mathbb{M})$ be non-empty at all? That is, why should a pmp action admit *any* point process factors? For a *free* action, this question is equivalent to constructing a cross section for the action $G \curvearrowright^\alpha (X, \mu)$. This can be done, but the proofs are too opaque for me to work with.

Recall that if $\Gamma \curvearrowright^\alpha (X, \mu) \preceq \Gamma \curvearrowright^\beta (Y, \nu)$, and α is free, then β must be free also (see Theorem 3.4 of [BK] for proof and further discussion).

Question 11. If Π weakly factors onto Υ, and Υ is free, then must Π be free also?

An alternative description of weak containment of actions $\Gamma \curvearrowright^\alpha (X, \mu)$, $\Gamma \curvearrowright^\beta (Y, \nu)$, at least for non-atomic spaces $(X, \mu), (Y, \nu)$, involves the polarising concept of *ultrafilters*. See [CKTD13] for the relevant definitions and the following theorem:

Theorem 37 (Conley, Kechris, Tucker-Drob)**.** Let $\mathcal{U}$ be a non-principal ultrafilter on $\mathbb{N}$, and $(X, \mu), (Y, \nu)$ non-atomic spaces. Then $\Gamma \curvearrowright^\alpha (X, \mu) \preceq \Gamma \curvearrowright^\beta (Y, \nu)$ if and only if the ultrapower of β factors onto α.

Work has been done by Carderi in [Car18] on ultraproducts of actions for lcsc groups, teasing out the various worrisome measurability complications that arise whenever the dreaded *ultras* arise.

Another failure was my inability to suss out *strong ergodicity* for point processes.

Definition 43. Let $G \curvearrowright (X, \mu)$ be a pmp action of an lcsc group. A sequence (A_n) of subsets of (X, μ) is *nontrivial* if $\mu(A_n)$ is uniformly bounded away from 0 and 1. It is *asymptotically invariant* if for all compact $K \subseteq G$ we have

$$\lim_{n \to \infty} \sup_{k \in K} \mu(A_n \triangle k A_n) = 0.$$

The action is said to be *strongly ergodic* if it has no nontrivial asymptotically invariant sequences.

Let $(Y, \nu, \mathcal{R})$ be a pmp cber. Its *full group* is

$$[\mathcal{R}] = \{T \in \mathrm{Aut}(Y, \nu) \mid (y, T(y)) \in \mathcal{R} \text{ for almost every } y \in Y\}.$$

A sequence of subsets (B_n) of (Y, ν) is said to be *nontrivial* if $\nu(B_n)$ is uniformly bounded away from 0 and 1. It is *asymptotically invariant* if for all $T \in [\mathcal{R}]$ we have

$$\lim_{n\to\infty} \nu(B_n \triangle T(B_n)) = 0.$$

The pmp cber is said to be *strongly ergodic* if it has no nontrivial asymptotically invariant sequences.

Question 12. Let μ be a free point process on G. Is $G \curvearrowright (\mathbb{M}, \mu)$ strongly ergodic if and only if the Palm equivalence relation $(\mathbb{M}_0, \mu_0, \mathcal{R})$ is strongly ergodic?

The most direct approach to proving a statement like this would be to construct a machine that converts asymptotically invariant sequences for the action into asymptotically invariant sequences for the equivalence relation and vice versa (it would also have to preserve nontriviality). There are ways to do this to go from a subset of $\mathbb{M}_0$ to one of $\mathbb{M}$, but it is the reverse direction that is unclear.

Bibliography

[ABB$^+$17] Miklos Abert, Nicolas Bergeron, Ian Biringer, Tsachik Gelander, Nikolay Nikolov, Jean Raimbault, and Iddo Samet. On the growth of L^2-invariants for sequences of lattices in Lie groups. *Ann. of Math. (2)*, 185(3):711–790, 2017.

[AD13] Claire Anantharaman-Delaroche. Invariant proper metrics on coset spaces. *Topology and its Applications*, 160(3):546 – 552, 2013.

[AN12] Miklós Abért and Nikolay Nikolov. Rank gradient, cost of groups and the rank versus Heegaard genus problem. *J. Eur. Math. Soc. (JEMS)*, 14(5):1657–1677, 2012.

[Avn05] Nir Avni. Spectral and mixing properties of actions of amenable groups. *Electronic Research Announcements of the American Mathematical Society*, 11(7):57–63, 2005.

[AW13] Miklós Abért and Benjamin Weiss. Bernoulli actions are weakly contained in any free action. *Ergodic theory and dynamical systems*, 33(2):323–333, 2013.

[BB99] Eric Babson and Itai Benjamini. Cut sets and normed cohomology with applications to percolation. *Proceedings of the American Mathematical Society*, 127(2):589–597, 1999.

[Bek18] Bachir Bekka. Spectral rigidity of group actions on homogeneous spaces. In *Handbook of group actions. Vol. IV*, volume 41 of *Adv. Lect. Math. (ALM)*, pages 563–622. Int. Press, Somerville, MA, 2018.

[BHI18] Lewis Bowen, Daniel Hoff, and Adrian Ioana. von Neumann's problem and extensions of non-amenable equivalence relations. *Groups Geom. Dyn.*, 12(2):399–448, 2018.

[BHM18] Francois Baccelli and Mir-Omid Haji-Mirsadeghi. Point-shift foliation of a point process. *Electron. J. Probab.*, 23:Paper No. 19, 25, 2018.

[BK] Peter J Burton and Alexander S Kechris. Weak containment of measure-preserving group actions. *Ergodic Theory and Dynamical Systems*, pages 1–53.

[Bla17] Bartlomiej Blaszczyszyn. Lecture Notes on Random Geometric Models — Random Graphs, Point Processes and Stochastic Geometry. Lecture, December 2017.

[BLPS99] Itai Benjamini, Russell Lyons, Yuval Peres, and Oded Schramm. Group-invariant percolation on graphs. *Geometric & Functional Analysis GAFA*, 9(1):29–66, 1999.

[Bow18] Lewis Bowen. All properly ergodic markov chains over a free group are orbit equivalent. *Unimodularity in Randomly Generated Graphs*, 719:155, 2018.

[Bow19] Lewis Bowen. Finitary random interlacements and the Gaboriau-Lyons problem. *Geom. Funct. Anal.*, 29(3):659–689, 2019.

[Car18] Alessandro Carderi. Asymptotic invariants of lattices in locally compact groups. *arXiv preprint arXiv:1812.02133*, 2018.

[CdlH16] Yves Cornulier and Pierre de la Harpe. *Metric geometry of locally compact groups*, volume 25 of *EMS Tracts in Mathematics*. European Mathematical Society (EMS), Zürich, 2016. Winner of the 2016 EMS Monograph Award.

[CKTD13] Clinton T. Conley, Alexander S. Kechris, and Robin D. Tucker-Drob. Ultraproducts of measure preserving actions and graph combinatorics. *Ergodic Theory Dynam. Systems*, 33(2):334–374, 2013.

[CLM18] A. Carderi and F. Le Maître. Orbit full groups for locally compact groups. *Trans. Amer. Math. Soc.*, 370(4):2321–2349, 2018.

[DVJ07] Daryl J Daley and David Vere-Jones. *An introduction to the theory of point processes: volume II: general theory and structure*. Springer Science & Business Media, 2007.

[For74] Peter Forrest. On the virtual groups defined by ergodic actions of R^n and $\mathbf{Z}^n$. *Advances in Math.*, 14:271–308, 1974.

[Fur99] Alex Furman. Orbit equivalence rigidity. *Annals of Mathematics*, 150(3):1083–1108, 1999.

[Fur09] Alex Furman. A survey of measured group theory. *arXiv preprint arXiv:0901.0678*, 2009.

[Gel11] Tsachik Gelander. Volume versus rank of lattices. *J. Reine Angew. Math.*, 661:237–248, 2011.

[GGP13] Ori Gurel-Gurevich and Ron Peled. Poisson thickening. *Israel J. Math.*, 196(1):215–234, 2013.

[GL09] Damien Gaboriau and Russell Lyons. A measurable-group-theoretic solution to von Neumann's problem. *Invent. Math.*, 177(3):533–540, 2009.

[Gre00] Schmidt Klaus Greschonig, Gernot. Ergodic decomposition of quasi-invariant probability measures. *Colloquium Mathematicae*, 84/85(2):495–514, 2000.

[HLS11] Alexander E. Holroyd, Russell Lyons, and Terry Soo. Poisson splitting by factors. *Ann. Probab.*, 39(5):1938–1982, 2011.

[HOU11] Cyril HOUDAYER. Invariant percolation and measured theory of nonamenable groups. *Séminaire BOURBAKI*, page 63, 2011.

[HP03] Alexander Holroyd and Yuval Peres. Trees and matchings from point processes. *Electron. Commun. Probab.*, 8:17–27, 2003.

[HP05] Alexander E. Holroyd and Yuval Peres. Extra heads and invariant allocations. *Ann. Probab.*, 33(1):31–52, 01 2005.

[Kec] Alexander S Kechris. The theory of countable borel equivalence relations.

[KEK05] Anthony W Knapp, Charles L Epstein, and Steven G Krantz. *Advanced real analysis*. Springer, 2005.

[Kin93] J. F. C. Kingman. *Poisson processes*, volume 3 of *Oxford Studies in Probability*. The Clarendon Press, Oxford University Press, New York, 1993. Oxford Science Publications.

[KL16] David Kerr and Hanfeng Li. Ergodic theory. *Springer Monographs in Mathematics. Springer, Cham*, 2016.

[KM04] Alexander S Kechris and Benjamin D Miller. *Topics in orbit equivalence*, volume 1852. Springer Science & Business Media, 2004.

[KPV15] David Kyed, Henrik Densing Petersen, and Stefaan Vaes. L^2-Betti numbers of locally compact groups and their cross section equivalence relations. *Trans. Amer. Math. Soc.*, 367(7):4917–4956, 2015.

[KST99] A.S Kechris, S Solecki, and S Todorcevic. Borel chromatic numbers. *Advances in Mathematics*, 141(1):1 – 44, 1999.

[LN11] Russell Lyons and Fedor Nazarov. Perfect matchings as IID factors on non-amenable groups. *European J. Combin.*, 32(7):1115–1125, 2011.

[LP11] Günter Last and Mathew D Penrose. Poisson process fock space representation, chaos expansion and covariance inequalities. *Probability Theory and Related Fields*, 150(3-4):663–690, 2011.

[LP16] Russell Lyons and Yuval Peres. *Probability on trees and networks*, volume 42 of *Cambridge Series in Statistical and Probabilistic Mathematics*. Cambridge University Press, New York, 2016.

[LP18] Günter Last and Mathew Penrose. *Lectures on the Poisson process*, volume 7 of *Institute of Mathematical Statistics Textbooks*. Cambridge University Press, Cambridge, 2018.

[LS99] Russell Lyons and Oded Schramm. Indistinguishability of percolation clusters. *Ann. Probab.*, 27(4):1809–1836, 10 1999.

[Mar12] Sébastien Martineau. Ergodicity and indistinguishability in percolation theory. *arXiv preprint arXiv:1210.1548*, 2012.

[MI17] James T Murphy III. Point-shifts of point processes on topological groups. *arXiv preprint arXiv:1704.08333*, 2017.

[nh] nullUser (https://mathoverflow.net/users/24840/nulluser). When does a stationary point process on group g have 0 or ∞ many points a.s.? MathOverflow. URL:https://mathoverflow.net/q/241933 (version: 2016-07-14).

[Paq18] Elliot Paquette. Distributional lattices on riemannian symmetric spaces. *Unimodularity in Randomly Generated Graphs*, 719:63, 2018.

[Pet13] Henrik Densing Petersen. L2-betti numbers of locally compact groups. *Comptes Rendus Mathematique*, 351(9-10):339–342, 2013.

[Sak14] J Sakai. *Settlers: The mythology of the white proletariat from mayflower to modern.* PM Press, 2014.

[Slu17] Konstantin Slutsky. Lebesgue orbit equivalence of multidimensional borel flows: A picturebook of tilings. *Ergodic Theory and Dynamical Systems*, 37(6):1966–1996, 2017.

[TD15] Robin D. Tucker-Drob. Weak equivalence and non-classifiability of measure preserving actions. *Ergodic Theory Dynam. Systems*, 35(1):293–336, 2015.

[Tho03] Simon Thomas. Superrigidity and countable borel equivalence relations. *Annals of pure and applied logic*, 120(1-3):237–262, 2003.

[Tim04] Adam Timar. Tree and grid factors of general point processes. *Electron. Commun. Probab.*, 9:53–59, 2004.

[VJ03] David Vere-Jones. *An Introduction to the Theory of Point Processes: Volume I: Elementary Theory and Methods.* Springer, 2003.

[Zim84] Robert J. Zimmer. *Ergodic theory and semisimple groups*, volume 81 of *Monographs in Mathematics.* Birkhäuser Verlag, Basel, 1984.