A Point-Process Framework for Local Limits of Random Graphs

Alec Boyadjiev

June 14, 2026

Abstract. We develop a point-process framework for local weak limits of labelled random graphs. The framework represents rooted graphs by rooted adjacency processes, point processes on product spaces that encode both vertices and edges while retaining the label-space structure from which the graph is induced. Our main theorem shows that, under weak convergence of rooted adjacency processes, weak convergence of the induced rooted graphs in the Benjamini-Schramm topology is equivalent to radial tightness, a compact-containment condition for rooted neighborhoods in the underlying label space. This reduces weak convergence of graphs to adjacency process convergence together with a sharp non-escape criterion. We then combine this construction with Palm theory to treat uniform rooting of finite graphs and obtain local weak convergence results for a broad class of spatial inhomogeneous random graphs under general conditions, with concrete results for mixed-binomial and Cox vertex processes with labels in $\mathbb{R}^d$ and $\mathbb{Q}_p^d$.

Contents

1 Introduction

In 1959, Erdős and Rényi published one of the first foundational papers on random graphs [12], initiating what became the modern probabilistic study of large networks. Their program was driven by two structural questions:

1. When a collection of vertices is connected at random, what types of graph properties emerge?
2. When such a collection gets increasingly large, how can expected behavior be understood?

In the decades since then, random graph theory has become a significant field, in no small part due to its great utility in modelling very large real world networks where connection between nodes or vertices is governed by some stochastic connection rule. Within the field, two main paradigms are commonly distinguished: dense graphs, where vertices have very many connections between them, and sparse graphs, where vertices have very few connections between them.

When studying the asymptotic behavior of dense graphs, one often asks questions about global density. For example, if three vertices are selected at random, what is the probability that they form a triangle? In dense graphs, this probability can remain nonzero even as the graph becomes very large, so it gives meaningful information about the overall graph structure. However, local questions can become less informative in this setting. For instance, if every vertex has many connections, then a randomly chosen vertex will generally be part of a triangle with a probability that tends to one as the graph gets larger.

Conversely, sparse graphs are better suited to questions about asymptotic local structure. If three vertices are selected at random in a sparse graph, the probability that they form a triangle often tends to zero, simply because there are too few edges compared to the number of possible triples. But one can still ask meaningful local questions, such as whether a randomly selected vertex is part of some triangle, via its local neighborhood. This probability may remain non-trivial, because it depends on the structure around that vertex rather than some global density of triangles.

The latter area is where this thesis focuses. In sparse graphs, asymptotic behavior is studied via a means of convergence known as *local weak convergence*, which captures the probability that the local neighborhood around a uniformly selected vertex "looks like" a given local neighborhood as the graph gets infinitely large. While this type of convergence is classically shown via combinatorial methods, such as in [17, Sec. 2.4.5], more recent work has begun to ask if this convergence can be understood via spatial representations of the graphs in question, where vertices are given labels in some space, often a metric space.

A perfect example of such an approach is the theory of *spatial inhomogeneous random graphs* as presented in [18], and the later work on point process representations of such models in [33]. The basic idea is that instead of viewing edges in a graph as connections between points, the edges themselves can be under-

stood as points in a product space. For example, if the vertices of the graphs are labelled in $\mathbb{R}$, one can view the edges as points in $\mathbb{R}^2$. As such, the edge between a vertex labelled 1.5 and a vertex labelled 2.7 would be represented by the point $(1.5, 2.7)$. Since the study of random collections of points has a well-developed means of analysis through point process theory, one can consider the question of lifting convergence of such representations in their associated spaces to convergence of the graphs that they represent in the local weak sense.

This thesis answers exactly the question of this lift. By abstracting and generalizing the proofs in [18] and related work, we develop a modular framework through which one can determine the convergence of a broad range of random graph models. We first define the specifics of our representation in Section 3, and then classify the exact criteria for such a lift of convergence for general rooted graphs via a single compact containment condition in Section 4. After that, Section 5 specializes the theory to rooted graphs where the root is chosen uniformly, as in local weak convergence, and provides an analogous theorem.

Section 6 then provides a very general class of random graph models, a variant of spatial inhomogeneous random graphs, on which more specific checkable criteria are established, and Section 7 then verifies the theory on concrete model families and examples. Ultimately, we validate the framework by generalizing results from [18] under marginally differing conditions, and demonstrate the utility of our framework on more complex vertex distributions and label spaces. For example, we determine local weak convergence for families of models on the p-adic fields $\mathbb{Q}_p$, and also demonstrate local weak convergence under Cox vertex processes, a highly general way of distributing vertices of space.

Ultimately, our framework is intended to serve as a modular baseline for determining local weak convergence of spatial random graph models, on which deeper theory and results can be built. By isolating and characterizing the conditions under which spatial convergence of graph representations can be lifted to local weak convergence of graphs, we hope to offer a common language for studying these models, so that new questions can be approached through an established structure rather than having to verify per-model local weak convergence from scratch. At a broader level, we hope that this perspective helps connect geometric intuition and probabilistic limit behavior in a way that makes the subject more coherent, more transferable, and easier to build on.

As a brief organizational note, we place most proposition and lemma proofs at the end of each section rather than immediately after their statements. This is intended to improve readability by allowing the main intuition of the argument to remain visible, while technical details are handled in one place. Unless stated otherwise, the proof of a given result appears in the propositions or lemmas subsection of the section in which it is stated. It is also helpful to keep Appendix C and Appendix D in mind, as they contain all of the notation and shorthand used repeatedly in the thesis. Appendix E may also be useful as it collects the individual cited results that we use. Some technical proofs and constructions are deferred to Appendix A and Appendix B respectively.

2 Preliminaries

Throughout this work, ideas from various fields will be tied together in order to build our theory. We therefore expect some level of familiarity from the reader, and refer to some sources where the terms we use can be found. The two primary areas that our work fits into are measure theory and graph theory. As the sections progress, we start using a few more specialized areas, namely Benjamini-Schramm theory and point process theory with a focus on Palm conditioning. It is also important for the reader to have a basic familiarity of topology, and a minor amount of group theory. It should be noted that most of the non-elementary constructions and results that we use are restated in text.

For basic measure theory, we expect familiarity with σ-algebras, measures, counting measures, measurable maps, basic probability theory with probability spaces and basic manipulation of probabilities and expectations, Borel σ-algebras and the Lebesgue measure, convergence in distribution (weak convergence), and a few mapping theorems (Portmanteau, continuous mapping theorem). For these we refer the reader to [22, 5, 19]. Building on this, we also expect some elementary understanding of point processes, but one should be able to get by if they sufficiently understand abstract random elements as measurable maps from probability spaces, as we explicitly state point process definitions and results in text. Perhaps a basic familiarity with the definition of a point process, along with the notion of a simple point process and some basic examples such as Cox and Poisson processes, would support the reading, and for this we refer the reader to [7, 25].

For basic graph theory, we expect familiarity with undirected simple graphs, graph distance, graph balls, neighbors, rooted graphs, isomorphism and rooted isomorphism classes, infinite graphs, local finiteness, and connected components. A few high-level references are [9, 6, 16]. We also refer the reader to [3] for introductory notions around the Benjamini-Schramm space of locally finite connected rooted graphs, though we also restate the key ideas in text.

A lot of our results are built on a topological foundation, and thus some elementary knowledge of topology is also beneficial. This includes the definition of a topology, subspace topologies, product topologies, open / closed sets, continuity, compactness, topological convergence, metric spaces, locally compact spaces, Polish spaces, the vague topology, and the Borel topology. We refer the reader to [29, 21, 5] for these notions.

On the topic of groups, it suffices to have a basic idea of the definition of a group, as well as group actions and measures / topologies on groups (via the Haar measure). For these one can look at [10]. Note however that these ideas will be restated in text. Lastly, for those curious, it may be helpful to understand some of the surrounding literature on random graphs, primarily [18] on spatial inhomogeneous random graphs, as much of our theory is an abstraction and modularization of the proof strategies used therein. However, this is not strictly necessary for the reading of this thesis.

With these preliminaries, one is well-equipped to understand our theory. We now take the time to give a few foundational definitions and results that are too central to defer to references. The first of these is the locally compact Polish space, or LC Polish space for short. Square LC Polish spaces are also defined here, which are necessary for our graph representations (adjacency processes) to behave properly. We use these as label spaces for vertices and edges. It should be noted that in most literature these spaces do not come with an attached metric, whereas in our work we will treat them as having an associated metric, as it is useful for our arguments. Practically, when using our theory, one would most often use label spaces like $\mathbb{R}$, $\mathbb{Q}_p$, or $\mathbb{Z}$ with their canonical metrics, or product spaces of them.

Definition 1 LC and Square LC Polish Spaces

Throughout, an *LC Polish space* S means a locally compact Polish space equipped with a fixed proper compatible metric d_S. In particular, bounded subsets of S are relatively compact, i.e. have compact closure.

A *square LC Polish space* S^2 means the product space $S \times S$, equipped with a fixed but arbitrary product metric induced by d_S. This product metric is again proper and compatible with the product topology.

Recall that a *Polish space* is a topological space that is separable and completely metrizable. A metric is *proper* if every closed bounded set (equivalently, every closed ball) is compact, which gives compactness control on bounded regions. A space is *second countable* if its topology has a countable basis. Every Polish space is second countable, which means that from a countable dense subset and rational radii, one obtains a countable basis of open balls. In our setting, this provides countable Borel generators and supports sequence-based measurability and approximation arguments. We now give the Portmanteau and continuous mapping results that we use on the corresponding Borel σ-algebras.

▶ **Proposition 2.1** Used Portmanteau Consequences. *Let (E, d) be a metric space, and let $(X_n)_{n\in\mathbb{N}}, X$ be E-valued random elements. Then, $X_n \Rightarrow X$ is equivalent to any of the following:*

1. *For all bounded and measurable*[1] *$f : E \to \mathbb{R}$, with X almost surely a continuity point of f (equivalently, $\mathbb{P}(X \in D_f) = 0$, where D_f is the discontinuity set of f), it holds that*

$$\mathbb{E}[f(X_n)] \to \mathbb{E}[f(X)].$$

2. *For every closed set $F \subseteq E$, it holds that*

$$\limsup_{n\to\infty} \mathbb{P}(X_n \in F) \leq \mathbb{P}(X \in F).$$

These are direct consequences of the Portmanteau theorem.

[1] To show $X_n \Rightarrow X$, testing against fully continuous functions suffices.

Here, $\Rightarrow$ denotes convergence in distribution. This result is used whenever we need to move from weak convergence to concrete probability and expectation statements. In practice, it is used to pass limits through bounded observables whose discontinuities are negligible at the limit, and to control bounds for events defined by closed sets.

▶ **Proposition 2.2** Continuous Mapping Theorem. *Let (E_1, d_1), (E_2, d_2) be metric spaces, let $\varphi : E_1 \to E_2$ be measurable, and let U_φ be its set of discontinuity points. If $X_n \Rightarrow X$ as E_1-valued random elements and $\mathbb{P}(X \in U_\varphi) = 0$, then*

$$\varphi(X_n) \Rightarrow \varphi(X).$$

This is the classical continuous mapping theorem.

This result is used whenever a convergence statement must be pushed through a map that turns one type of random object into another. In practice, we first verify continuity at the relevant limit points, and then use this theorem to conclude convergence after applying the map.

3 Adjacency Processes

The notion of a point process representation of a graph is not unfamiliar in random graph theory, being present in the literature for over a decade (for a foundational source, see the use of *adjacency measure* in Veitch et al. [36]). In addition to the original context as representations of Graphexes (under exchangeability conditions), more modern work including van der Hofstad et al. [18] and Sánchez Flores [33] has framed such point processes as representations of the random rooted graphs resulting from local weak limits in sparse regimes, where Palm distributions are used to provide a fixed root from which to explore the distribution of local neighborhoods.

Our work provides a study and generalization of the latter, reformulating the proof strategies of [18, 33] as general machinery that can be applied to a broader range of vertex distributions under suitable conditions on the connection kernel that generates the graphs. This first section will state and motivate our definition of the adjacency measure, which varies slightly from the conventions of the literature, and then establish some general results on them that will be used throughout the paper. We will treat the adjacency measures as the fundamental objects, from which graphs are constructed, rather than constructing them retroactively to fit an existing graph model as in [33].

3.1 Graphs as Their Edges

When looking at the definition of an undirected simple graph $G = (V, E)$, it is clear that essentially all of the "useful" information lives in the edges. More specifically, with the exception of lone vertices, all of the vertex information of the graph can be recovered from its edges. Thus, when working with labels in some LC Polish space S, a graph (up to lone vertices) can be viewed as a collection of points in S^2. Since we want to analyze random versions of these representations, the most sensible option is to view them as random collections of points. These are called point processes, defined as random counting measures. Thus, we define our representation using counting measures instead of sets. A counting measure is just a measure that is integer-valued.

This is exactly the approach currently used in the literature. For any simple graph G with vertex labels in S, the corresponding[2] Veitch-Roy adjacency measure is the simple symmetric counting measure counting each of the edges in S^2 [36, Def. 4.1], where symmetry is used to mark both orderings of the undirected edges. Since lone vertices are typically irrelevant in these theories, they can be ignored by the representation. Here, symmetric means that whenever (x, y) appears, so does (y, x). When such a measure is taken randomly, it can simply be called an *adjacency process*.

In this work, we establish weak convergence of graphs constructed from a similar

[2] Strictly speaking, these measures must be boundedly finite, so not all labellings are representable; we ignore this here for clarity of motivation but impose it formally later.

class of adjacency processes by deriving sufficient conditions under which convergence of the underlying spatial processes propagates to convergence in the Benjamini–Schramm sense. The primary issue to solve is to translate "closeness" of the points in an adjacency measure into "closeness" of the isomorphism classes of the constructed graphs. For this goal, Veitch-Roy adjacency measures do not suffice, as even small perturbations of a Veitch–Roy adjacency measure can change the isomorphism class of the constructed graph, even when the measures are arbitrarily close in the natural point process topology. Consider the following diagram:

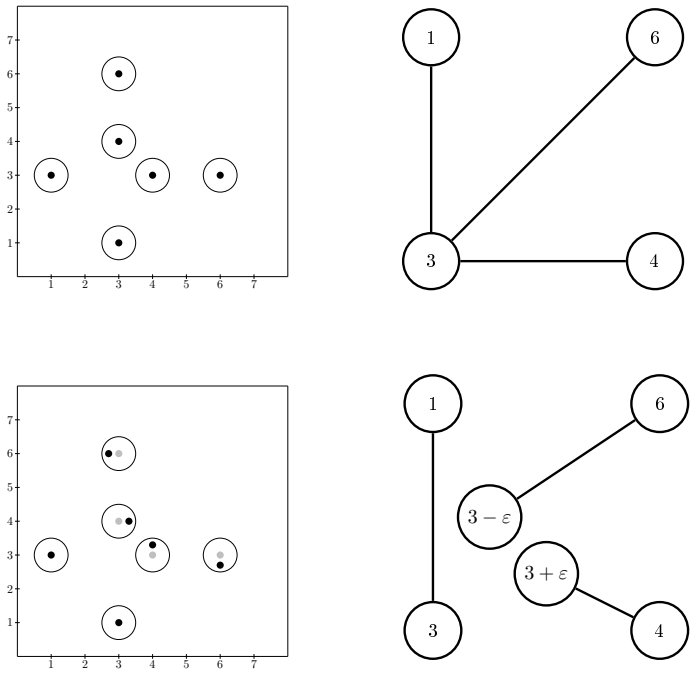

Figure 1: An arbitrarily small change breaks graph equivalence.

In the figure, a small perturbation of two atoms (and their reflections) in S^2 results in edges being reassigned to different endpoints. Despite the perturbation being arbitrarily small in the ambient space, the induced graph is no longer isomorphic to the original. This reflects a structural issue: the representation encodes edges, but does not encode vertex identity in a way that is stable under perturbations. Consequently, spatial proximity of adjacency measures does not translate to combinatorial proximity of the induced graphs.

This instability arises because vertex locations are not anchored to the edges that define them. To address this, we modify the representation by imposing a consistency condition: whenever an edge (x, y) is present, the diagonal points (x, x) and (y, y) must also be included. This can be interpreted as anchoring each

vertex within the measure, thereby coupling vertex identity to edge structure. While this modification may appear minor, it fundamentally alters the behavior of the representation, and removes bad cases like the one in Figure 1. Its precise role in ensuring stability and enabling convergence results will be formalized in Section 4.2. Additionally, including the diagonal allows us to condition the adjacency structure on a vertex (in a Palm sense) instead of just on an edge, which is fundamental to our later results, and would not be possible via Veitch-Roy adjacency measures. The exact use of Palm distributions can be seen in Section 5.1. With motivation complete, we give the formal definition of adjacency measure below.

The definition will use the term boundedly finite. In some work, the term locally finite is used. However, this is problematic in our work, as locally finite is already a term used to characterize infinite graphs where each vertex has finite degree. Thus we distinguish them, and give an alternative spatial condition that corresponds to local finiteness of the corresponding graph. This is formalized in Section 3.2, when we discuss graph construction. For clarity, we restate the definition of boundedly finite below.

Definition 2 Boundedly Finite Measure

A measure μ on an LC Polish space[a] S is *boundedly finite* iff for every bounded Borel set $B \subseteq S$, $\mu(B) < \infty$.

[a] S^2 as in Def. 3 is also LC Polish.

With this, we can define an adjacency measure and the class of such measures.

Definition 3 Adjacency Measure

An *adjacency measure* is a boundedly finite simple counting measure τ on a square LC Polish space S^2 with the property that for all (x, y) with $\tau(\{(x, y)\}) = 1$, it also holds that

$$\tau(\{(x, x)\}) = 1, \qquad \tau(\{(y, y)\}) = 1.$$

The space of adjacency measures on S^2 is denoted $\mathcal{A}(S^2) \subset \mathcal{N}(S^2)$, where $\mathcal{N}(S^2)$ is the space[a] of all boundedly finite simple counting measures on S^2. We equip $\mathcal{A}(S^2)$ with the subspace topology inherited from $\mathcal{N}(S^2)$ under the vague topology [30, p. 28]. This is the topology in which counting measures converge if and only if their integrals against all continuous compactly supported test functions converge.

[a] Often denoted $\mathcal{N}_{bf}$; we write $\mathcal{N}$ as we do not need the more general space.

Henceforth we will use adjacency measure to refer to the above, and Veitch-Roy adjacency measure to refer to the reference if necessary. We will use τ always

and exclusively to denote adjacency measures, while μ and ν remain available for general measures.

It should be noted that in the above definition, symmetry has been removed. This is a design choice. We have chosen to let the graph construction handle the removal of diagonal points and the inclusion of an edge if either one of its orderings is present in the adjacency measure, to allow more flexibility in models and arguments. However, this is a strictly broader class than symmetric adjacency measures, and thus our results can also be applied in the symmetric case whenever necessary without alteration.

Thus, our adjacency measures should not be read as literal edge representations. They are regularized representations carrying both vertex and edge information. Conversely, the usual Veitch–Roy adjacency measure associated with the same graph is recovered[3] by removing the diagonal and symmetrizing. Hence the diagonal does not introduce self-edges, it only changes the representation from which the simple graph is constructed. Our adjacency measures also allow for lone vertices to be represented, but this is minor and usually inconsequential.

There are a few more things to define before we move on to graph construction. Perhaps the most fundamental is that our space of graphs is the Benjamini–Schramm space, which only considers rooted locally finite graphs. Thus the natural way to construct a graph from an adjacency measure is to have a root associated with it.

Definition 4 Rooted Adjacency Measure

An *rooted adjacency measure* is a pair (τ, x) where τ is an adjacency measure on a square LC Polish space S^2 and $x \in S$ satisfies

$$\tau(\{(x, x)\}) = 1.$$

The space of rooted adjacency measures on S^2 is denoted $\mathcal{A}^\circ(S^2) \subset \mathcal{N}(S^2) \times S$. We equip $\mathcal{A}^\circ(S^2)$ with the subspace topology inherited from $\mathcal{N}(S^2) \times S$ under the product of vague topology and Borel topology on S.

We must also formally define the (rooted) adjacency process. This is simply a random variable taking values in one of the above spaces.

Definition 5 (Rooted) Adjacency Process

An *adjacency process* ξ on a square LC Polish space S^2 is a random variable taking values in $\mathcal{A}(S^2)$. Similarly, a *rooted adjacency process* (ξ, p) is a random variable taking values in $\mathcal{A}^\circ(S^2)$.

Thus an adjacency process is a type of point process, while a rooted adjacency process is a point process with an associated random root. We use ξ always and

[3] Technically only recoverable in the case where there are no lone vertices.

exclusively to denote adjacency processes, and reserve η for other point processes. For completeness, we define the spaces of symmetric adjacency measures as well.

Definition 6 Symmetric Adjacency Measure

An adjacency measure τ on a square LC Polish space S^2 is *symmetric* if for every $(x, y) \in \tau$, it also holds that $(y, x) \in \tau$. The space of symmetric adjacency measures is denoted

$$\mathcal{A}_{\text{sym}}(S^2) := \{\tau \in \mathcal{A}(S^2) : \tau \text{ is symmetric}\}.$$

The rooted symmetric adjacency-measure space is denoted

$$\mathcal{A}^{\circ}_{\text{sym}}(S^2) := \mathcal{A}^{\circ}(S^2) \cap \big(\mathcal{A}_{\text{sym}}(S^2) \times S\big).$$

Both spaces inherit their respective subspace topologies from $\mathcal{A}(S^2)$ and $\mathcal{A}^{\circ}(S^2)$. We also give the map $\text{Sym} : \mathcal{A}(S^2) \to \mathcal{A}_{\text{sym}}(S^2)$ by

$$\text{Sym}(\tau) := \tau + \sum_{(x,y)\in\tau} 1_{\{(y,x)\notin\tau\}}\delta_{(y,x)},$$

which converts adjacency measures into symmetric variants. We overload the notation to also have $\text{Sym}(\tau, x) := (\text{Sym}(\tau), x)$ for rooted adjacency measures.

Here the shorthand $(x, x) \in \tau$ is used to denote $\tau(\{(x, x)\}) = 1$, and is used throughout the paper. In general, we will extend this convention to any set operations, and allow using simple counting measures in set operations as shorthand for the set of positive-measure singletons under the counting measure. See Appendix D for detail. The notation δ_z denotes the counting measure $\delta_z(A) := \mathbf{1}_{z\in A}$.

Lastly, it is useful to define the vertex map.

Definition 7 Vertex Map

Let τ be an adjacency measure on a square LC Polish space S^2. Then, the *vertex map* $V : \mathcal{A}(S^2) \to \mathcal{N}(S)$ maps

$$\tau \mapsto \sum_{(x,x)\in\tau} \delta_x.$$

This map is measurable, and thus can be applied to adjacency processes to give a *vertex process*. This is just a boundedly finite point process on S, but it is often convenient to refer to the "vertex process corresponding to an edge process", and thus we sometimes use the name even if the class has no distinguishing properties.

▶ **Lemma 3.1** Measurability of Vertex Map. *The vertex map V in Definition 7 is measurable.*

The proof is standard, and thus deferred to Appendix A.1. We also prove that all of the subspaces we have defined are LC Polish.

▶ **Lemma 3.2.** *For any square LC Polish S^2, the spaces*

1. $\mathcal{A}(S^2)$,
2. $\mathcal{A}^\circ(S^2)$
3. $\mathcal{A}_{sym}(S^2)$,
4. $\mathcal{A}^\circ_{sym}(S^2)$,

are all Polish under their respective subspace topologies, and the map Sym from Definition 6 is measurable under the respective Borel σ-algebras, and the map $(\tau, x) \mapsto (Sym(\tau), x)$ is measurable and well-defined from $\mathcal{A}^\circ(S^2)$ to $\mathcal{A}^\circ_{Sym}(S^2)$.

Again, the proof is standard and deferred to Appendix A.2. With the initial structure concluded, we can move on to the construction of graphs.

3.2 Construction of Graphs

Graph construction has some caveats. An adjacency measure naturally corresponds to a full graph, which may not be locally finite, may be empty, or may have disconnected components. However, the Benjamini-Schramm [3] space $\mathcal{G}_*$, which we use as our primary convergence space, records and compares rooted locally finite graphs through the local neighborhoods of the root. Consequently, once a root is chosen, only the connected component of that root is visible to the local topology. Concretely, $\mathcal{G}_*$ is the space of rooted isomorphism classes of connected, locally finite rooted graphs, equipped with the metric

$$d_{\mathcal{G}^*}(G, H) = 2^{-r(G,H)},$$

where $r(G, H)$ is the largest radius for which the rooted balls of the graphs are isomorphic, with the convention $d_{\mathcal{G}^*} = 0$ when all rooted balls agree. Graph isomorphism, as used above, is equality of graph structure up to relabelling of vertices, which in the rooted case must also preserve the distinguished root. Local finiteness is the condition that guarantees that the degree of any vertex is finite. Rooted balls are simply the finite neighborhoods around the distinguished root, containing all vertices within a fixed graph-distance radius from the root together with the edges between them, based on the graph distance $d_G(x, y)$, which gives the length of the shortest path from x to y in G.

Thus, we do a two stage graph construction. First, we construct the corresponding labelled graph object, and then construct the graphs in the appropriate spaces via graph isomorphism. First the former:

Definition 8 Constructed Labelled Graph (Undirected)

Let $\tau \in \mathcal{A}(S^2)$ be an adjacency measure in a square LC Polish space S^2. Then, the *constructed labelled graph* $G_S(\tau)$ is given by $G_S(\tau) = (V_\tau, E_\tau)$ where

$$V_\tau = \{y \mid (y, y) \in \tau\},$$

$$E_\tau = \{\{y, z\} \mid y \neq z,\ (y, z) \in \tau \text{ or } (z, y) \in \tau\}.$$

Similarly, if we consider a rooted adjacency measure (τ, x) for some $x \in \tau$, the *constructed labelled graph* $G_S(\tau, x)$ is simply the graph $(G_S(\tau), x)$.

The above gives us labelled graphs. From these, we can then construct our unlabelled graphs.

Definition 9 Constructed Graph (Undirected)

Let $(\tau, x) \in \mathcal{A}^\circ(S^2)$ be a rooted adjacency measure in a square LC Polish space S^2. Then, define the *constructed graphs*

$$G(\tau, x) = [\mathcal{C}_{G_S(\tau)}(x), x],$$

$$G(\tau) = [G_S(\tau)].$$

Where $\mathcal{C}_{G_S(\tau,x)}(x)$ denotes the connected component of x in $G_S(\tau, x)$, and the brackets denote graph isomorphism and rooted graph isomorphism classes respectively[a].

[a] Technically, the graph isomorphism is taken within some universe. Canonically, this is the space of all graphs with labels in $\mathbb{N}$.

As seen above, in the rooted case the construction retains only the connected component of the root. This is deliberate, as Benjamini–Schramm convergence is determined entirely by finite rooted balls, and therefore only sees the component reachable from the distinguished root. Components disconnected from the root are invisible to this local topology.

One could instead first construct a full rooted graph, possibly disconnected, and then define Benjamini–Schramm convergence on that larger class through the usual rooted-ball pseudometric[4]. However, this would require introducing a separate topology and measurable structure on the space of all rooted locally finite graphs, while the resulting local convergence would still depend only on the root component. We avoid this unnecessary intermediate structure by building the localization directly into the graph construction map. Thus the rooted graph construction takes values immediately in the standard Benjamini–Schramm space $\mathcal{G}_*$, and we maintain the notation G as it is the simplest for our

[4] Only a pseudometric on this larger class since distance is zero between graphs with differing non-root components.

context, even though we are not performing a full transparent graph construction like in the labelled variants.

We now give the property that guarantees local finiteness. Immediately after, we will prove the equivalence of this property to local finiteness of the constructed graph.

Definition 10 Fiber-Finite Adjacency Measure

Let $\tau \in \mathcal{A}(S^2)$ be an adjacency measure on a square LC Polish space S^2. We call τ *fiber-finite* if it has a finite count on every fiber. Specifically, this means that for every $x \in S$, it holds that

$$\tau(\{(x, y) \mid y \in S\}) < \infty,$$

$$\tau(\{(y, x) \mid y \in S\}) < \infty.$$

Intuitively, one can understand that these fibers represent incident edges to some vertex, and thus by keeping the counts finite, we keep the degrees finite. This is identical to the condition that the constructed graph is locally finite. The proof is left in-text to illustrate the argument.

▶ **Lemma 3.3.** *Let $\tau \in \mathcal{A}(S^2)$ be an adjacency measure in a square LC Polish space S^2. Then $G(\tau)$ is locally finite if and only if τ is fiber-finite in the sense of Definition 10.*

Proof. First note that the fibers are Borel subsets of S^2. Indeed, since S is Hausdorff, $\{x\}$ is closed for every $x \in S$, and hence both $\{x\} \times S$ and $S \times \{x\}$ are Borel. Let

$$G_S(\tau) = (V_\tau, E_\tau)$$

be the constructed labelled graph, so that

$$G(\tau) = [G_S(\tau)].$$

Since local finiteness is invariant under graph isomorphism, it suffices to work with the labelled representative $G_S(\tau)$. For $x \in S$, write

$$R_x(\tau) := \{y \in S : (x, y) \in \tau\}, \qquad C_x(\tau) := \{y \in S : (y, x) \in \tau\}.$$

Since τ is simple, the fiber counts are precisely

$$\tau(\{x\} \times S) = |R_x(\tau)|, \qquad \tau(S \times \{x\}) = |C_x(\tau)|.$$

Thus τ is fiber-finite if and only if

$$|R_x(\tau)| < \infty \quad \text{and} \quad |C_x(\tau)| < \infty \qquad \text{for every } x \in S.$$

($\Leftarrow$) Assume that τ is fiber-finite. Let $x \in V_\tau$. Every neighbor y of x in $G_S(\tau)$ satisfies $y \neq x$ and either $(x, y) \in \tau$ or $(y, x) \in \tau$. Hence, letting $N_{G_S(\tau)}(x)$ denote the graph neighbors of x,

$$N_{G_S(\tau)}(x) \subseteq R_x(\tau) \cup C_x(\tau).$$

By fiber-finiteness, both $R_x(\tau)$ and $C_x(\tau)$ are finite. Hence, $N_{G_S(\tau)}(x)$ is finite. Since this holds for every $x \in V_\tau$, the labelled graph $G_S(\tau)$ is locally finite. Therefore,

$$G(\tau) = [G_S(\tau)]$$

is locally finite.

($\Rightarrow$) Assume that $G(\tau)$ is locally finite. Then the representative $G_S(\tau)$ is locally finite. Let $x \in S$.

If $x \notin V_\tau$, then $(x, x) \notin \tau$. Since τ is an adjacency measure, any atom of the form (x, y) or (y, x) would imply $(x, x) \in \tau$. Therefore,

$$R_x(\tau) = C_x(\tau) = \varnothing,$$

and hence

$$|R_x(\tau)| = |C_x(\tau)| = 0 < \infty.$$

Now suppose $x \in V_\tau$. If $y \in R_x(\tau)$ and $y \neq x$, then $(x, y) \in \tau$, so $\{x, y\} \in E_\tau$. Thus

$$R_x(\tau) \subseteq \{x\} \cup N_{G_S(\tau)}(x).$$

Similarly,

$$C_x(\tau) \subseteq \{x\} \cup N_{G_S(\tau)}(x).$$

Since $G_S(\tau)$ is locally finite, the neighbour set $N_{G_S(\tau)}(x)$ is finite. Therefore both $R_x(\tau)$ and $C_x(\tau)$ are finite. Since $x \in S$ was arbitrary, τ is fiber-finite. ◀

The definition and lemma transfer also to rooted adjacency measures and rooted graphs, as the choice of root does not affect local finiteness. However, the implication becomes one way, as the rooted local construction only requires the connected component of the root to be locally finite, while fiber-finiteness is stronger: it makes the full constructed graph locally finite. We impose this global condition because our mechanisms for choosing roots, introduced in Section 3.3, are defined at the level of the full adjacency measure and may select any vertex. Thus, regularity guarantees that every admissible root gives an element of $\mathcal{G}_*$. We give this lemma formally below with a brief proof.

▶ **Lemma 3.4.** *Let $(\tau, x) \in \mathcal{A}^\circ(S^2)$ be a rooted adjacency measure in a square LC Polish space S^2. Then $G(\tau, x)$ is locally finite if τ is fiber-finite in the sense of Definition 10.*

Proof. Assume that τ is fiber-finite. By Lemma 3.3, the full constructed graph $G_S(\tau)$ is locally finite. The rooted graph $G(\tau, x)$ is the rooted isomorphism class of the connected component of x in $G_S(\tau)$. Since every connected component of a locally finite graph is locally finite, the component of x is locally finite. Hence $G(\tau, x)$ is locally finite. ◀

For the later results, it is important to define the subspace of admissible adjacency measures, i.e. those that map measurably into $\mathcal{G}_*$. To do that, we have the following definition and lemma.

Definition 11 Regular Adjacency Measures

An adjacency measure[a] is called *regular* if it is non-empty and fiber-finite. We denote the space of all regular adjacency measures by

$$\mathcal{E}(S^2) \subset \mathcal{A}(S^2),$$

equipped with the subspace topology inherited from $\mathcal{A}(S^2)$. Similarly, we denote the space of all regular rooted adjacency measures by

$$\mathcal{E}^\circ(S^2) \subset \mathcal{A}^\circ(S^2),$$

again equipped with the subspace topology.

[a] An adjacency measure in the sense of Def. 3, not in the Veitch-Roy sense [36].

The term regular encapsulating these properties is used for convenience, to keep theorem statements concise. Finally, we can give the lemma that makes the graph construction usable.

▶ **Lemma 3.5** Measurability of Graph Construction. *For any square LC Polish space S^2, the map $G : \mathcal{E}^\circ(S^2) \to \mathcal{G}_*$ given by $(\tau, x) \mapsto G(\tau, x)$ is well-defined and measurable from the Borel σ-algebra on $\mathcal{E}^\circ(S^2)$ to the canonical [5] σ-algebra on $\mathcal{G}_*$.*

The proof is technical, relying on measurable atom relations, measurable finite set-valued selections, and a canonical enumeration of the rooted connected component. It is therefore deferred to Appendix A.3.

This measurability is key to our analysis, as we can now apply graph construction to not only regular adjacency measures, but also regular adjacency processes. Furthermore, we define a collection of additional maps that we will sometimes use, and prove that all of these are measurable.

[5] This is also the Borel σ-algebra emerging from the Benjamini-Schramm metric, see [3].

Definition 12 Graph Maps

Define the following maps:

- $B_r : \mathcal{G}_* \to \mathcal{G}_*$ maps $[G, x]$ to the rooted graph ball $[B_r(G, x), x]$ of radius r.
- $V_r : \mathcal{E}^\circ(S^2) \to \mathcal{N}(S)$ is defined by

$$V_r(\tau, x) := \sum_{y \in W_r(\tau, x)} \delta_y,$$

where

$$W_r(\tau, x) := \{y \in V_\tau : d_{G_S(\tau)}(x, y) \leq r\}.$$

The rooted ball is self-explanatory, but the map V_r can be essentially seen as the collection of vertices that lead to the rooted ball at radius r. Of course, we also need to prove that these are well-defined and measurable.

▶ **Lemma 3.6.** *For any square LC Polish space S^2, the maps from Definition 12 are well-defined and measurable on the respective σ-algebras* [6].

The proof is standard and deferred to Appendix A.4. We also introduce the shorthand $B_r(\tau, x) := B_r(G(\tau, x))$ which is also measurable under composition. Before we can continue with our results, we need one final bit of structure. We need to, when given a regular adjacency measure, be able to select a root according to some distribution. For this, we use the theory of Markov kernels.

3.3 Rooting Kernels

A Markov kernel[7] is a measurable assignment of a probability measure on one space to each element of another. For a reference, see Klenke [22, Def. 8.25]. In our context, our elements are regular adjacency measures on a square LC Polish space S^2, and our probability measures are over the space S, in such a way that only points within the adjacency measure have positive selection probability. We call such a Markov kernel a *rooting kernel*.

Definition 13 Rooting Kernel

A *rooting kernel* acting on regular adjacency measures in a square LC Polish space S^2 is a Markov kernel $R : \mathcal{L} \subseteq \mathcal{E}(S^2) \times \mathcal{B}(S) \to [0, 1]$ such that

$$R(\tau, S \setminus V(\tau)) = 0.$$

The above gives a distribution of possible roots, selected from the vertices of the regular adjacency measure. Regularity is only required here to avoid empty

[6] Each Borel with respect to the relevant topology.
[7] Also sometimes called a transition or probability kernel.

measures, but it is likely that this definition can be extended to all non-empty adjacency processes since local finiteness is not the constraining factor.

In the definition, we only required the kernel to be defined on a subspace $\mathcal{L} \subseteq \mathcal{E}(S^2)$. This is intentional. Often, one wants a rooting kernel that is natural on a restricted class of adjacency measures but has no sensible extension to all of $\mathcal{E}(S^2)$. The main example is uniform rooting: it is only well-defined for finite (where finite also means non-empty) adjacency measures. Allowing subspaces is therefore essential, since uniform rooting is central to our analysis. Below, we give the formal definition, and then prove that this is a kernel.

Definition 14 Uniform Rooting Kernel

Let S^2 be a square LC Polish space. Define the *uniform rooting kernel* $\mathbb{U}$ on the subspace $\mathcal{L} \subseteq \mathcal{E}(S^2)$ of finite[a] adjacency measures ($0 < \tau(S^2) < \infty$) as

$$\mathbb{U}(\tau, A) := \frac{|V(\tau) \cap A|}{|V(\tau)|}.$$

[a] Finite will always imply non-empty when used for adjacency measures.

The above is indeed a rooting kernel.

▶ **Lemma 3.7.** *For any square LC Polish space S^2, $\mathbb{U}$ is a rooting kernel.*

This proof is also standard, deferred to Appendix A.5. While many other interesting rooting kernels exist, the uniform rooting kernel is the only one that we will use within the scope of this paper, so we leave out any other examples. Thus, we have finally concluded the core structures used in the paper, and we move on to the stating and the proofs of our results, starting with the most general but least applicable theorem for weak convergence of graphs constructed by regular rooted adjacency processes.

4 Weak Graph Convergence

The choice of the adjacency measure representation brings up a natural question: If one has a sequence of regular rooted adjacency processes $(\xi_n, p_n)_{n\in\mathbb{N}}$ on a square LC Polish space S^2 and some limiting regular rooted adjacency process $(\xi_\infty, p_\infty) \in \mathcal{E}^\circ(S^2)$, satisfying

$$(\xi_n, p_n) \Rightarrow (\xi_\infty, p_\infty)$$

in $\mathcal{E}^\circ(S^2)$, when does it hold that

$$G(\xi_n, p_n) \Rightarrow G(\xi_\infty, p_\infty)$$

in $\mathcal{G}_*$ as well? Here, $\Rightarrow$ denotes convergence in distribution (weak convergence) of the random objects in the respective topologies.

This is in some sense an abstraction of the classical *local weak convergence*, defined to be the weak convergence in $\mathcal{G}_*$ of a sequence of finite unrooted graphs under a uniform root selection. However, despite this statement being more general, it has a very clean sufficient condition, which reduces analysis of the specialized local weak case to showing that this condition holds. This is exactly the approach that we take.

It is very important to clarify that convergence of these regular rooted adjacency processes in $\mathcal{N}(S^2)$ is not sufficient. The convergence must occur within $\mathcal{E}^\circ(S^2)$, as otherwise the graph construction operation is not well-defined for the limit object.

In order to motivate our sufficient condition, we will give an example to illustrate the core mechanism by which this implication may not hold. The intuition behind it is that the graph construction process is indifferent to distances between vertices, whereas the spatial representation is not. Specifically, the sequence could have a subset of its points (representing edges/vertices) tend to infinity and thus disappear from the limiting representation. However, the corresponding sequence of graphs would not capture this change, and thus would reach a different limit.

To see this, consider a sequence of deterministic embeddings into $\mathbb{N}$ of K_2, the connected graph on two vertices. We label one vertex with 0, and one with n, and represent the edge in the adjacency measure by both $(0, n)$ and $(n, 0)$. We then fix the root to be at 0. Formally, define for $n \in \mathbb{N}$ the rooted adjacency process

$$\xi_n := \delta_{(0,0)} + \delta_{(n,n)} + \delta_{(0,n)} + \delta_{(n,0)}, \qquad p_n := 0.$$

These are all evidently regular rooted adjacency processes[8], and the sequence $(\xi_n, p_n)_{n\in\mathbb{N}}$ converges weakly in $\mathcal{E}^\circ(S^2)$ to $(\xi_\infty, p_\infty) := (\delta_{(0,0)}, 0)$, since all other points vanish as $n \to \infty$.

[8] We view the deterministic case as a random variable with a single full probability outcome.

In $\mathcal{G}_*$, however, we find that for all $n \in \mathbb{N}$, the corresponding graph $G_n := G(\xi_n, p_n)$ is exactly K_2 with the root placed at one of the two vertices (due to symmetry, the root choice is irrelevant as the rooted graph isomorphism class is identical from all roots). Thus, the graph sequence is a sequence that is always equal, and it converges weakly to $G_\infty := K_2$.

However, the constructed graph $G(\xi_\infty, p_\infty)$ is instead K_1, the graph containing a single vertex. Thus $G_\infty \neq^d G(\xi_\infty, p_\infty)$, where $=^d$ denotes equality in distribution, and the implication does not hold.

The necessary fix is clear: We need to, in some way, constrain the points within the sequence of adjacency measures. But how should this be done? One could, for example, require all the adjacency processes to almost surely be contained within a compact set. However, this is far too restrictive of a condition, and prevents useful limiting behavior. Instead, the most general condition is a condition that we refer to as *radial tightness*. This should not be confused with Prokhorov style tightness [22, Def. 13.26]. Instead, this condition simply ensures that for every test radius r we can find some fixed compact set L such that the cloud of vertex labels corresponding to the rooted ball of radius r is eventually contained within L with arbitrarily high probability. Specifically,

Definition 15 Radial Tightness

A sequence $(\xi_n, p_n)_{n\in\mathbb{N}}$ taking values in $\mathcal{E}^\circ(S^2)$ is called *radially tight* iff for every $r \in \mathbb{N}$ and every $\varepsilon > 0$, there exist a compact set $L \subset S$ and an index $N \in \mathbb{N}$ such that for all $n \geq N$

$$\mathbb{P}(V_r(\xi_n, p_n) \subset L) \geq 1 - \varepsilon.$$

Equivalently,

$$\mathbb{P}(V_r(\xi_n, p_n) \not\subset L) < \varepsilon.$$

This is a very general condition, and in fact, it is actually equivalent to graph convergence once the convergence of the adjacency processes has been assumed. With the setup complete, we can finally give the general weak convergence theorem.

4.1 Weak Convergence Theorem

We will first state the weak convergence theorem, and then give the proof up to a collection of propositions and lemmas.

▶ **Theorem 4.1** Weak Convergence Theorem. *Let $(\xi_n, p_n)_{n\in\mathbb{N}}$ be a sequence of regular rooted adjacency processes on a square LC Polish space S^2. Suppose there exists a limiting regular rooted adjacency process (ξ_∞, p_∞) on S^2 for which*

$$(\xi_n, p_n) \Rightarrow (\xi_\infty, p_\infty).$$

Then, it holds that

$$G(\xi_n, p_n) \Rightarrow G(\xi_\infty, p_\infty)$$

if and only if $(\xi_n, p_n)_{n\in\mathbb{N}}$ *is radially tight.*

The value of this result is that it separates the spatial and graph-theoretic parts of the argument. Once weak convergence of the underlying adjacency processes has been established, the corresponding graph convergence is equivalent to a single compact-containment condition, rather than requiring model-specific arguments about the induced graphs themselves. This splits graph convergence questions into two conceptually distinct tasks. First, prove convergence of the spatial representations, and then verify radial tightness of the rooted neighborhoods. In the remainder of the paper, we will break down these conditions, eventually resulting in concrete results for various model families.

The intuition behind this theorem carries over cleanly from the prior exposition: Preventing escaping vertices leads to good behavior. In order to present the proof, we first need a handful of lemmas and propositions. These will be stated below, but their proofs will be deferred to later in this section for clarity. Firstly, we need to verify that any chosen compact set L satisfying radial tightness also contains the rooted ball of the limit with high probability. Specifically,

▶ **Lemma 4.2.** *Let* $(\xi_n, p_n)_{n\in\mathbb{N}}$ *be a sequence of regular rooted adjacency processes on a square LC Polish space* S^2 *converging to some limit regular rooted adjacency process* (ξ_∞, p_∞) *on* S^2*. Then, for any* $r \in \mathbb{N}$ *and* $\varepsilon > 0$*, if there exist a compact set* $L \subset S$ *and an index* $N \in \mathbb{N}$ *such that for all* $n \geq N$*,*

$$\mathbb{P}(V_r(\xi_n, p_n) \subset L) \geq 1 - \varepsilon,$$

then it holds that

$$\mathbb{P}(V_r(\xi_\infty, p_\infty) \subset L) \geq 1 - \varepsilon,$$

as well.

In addition to this, we also need to verify that for every compact set in S and regular adjacency process on S^2, we can find a larger compact set in S with good boundary behavior. The exact statement is given below.

▶ **Lemma 4.3.** *Let* ξ *be a regular adjacency process on a square LC Polish space* S^2*. Then, for every compact set* $\tilde{L} \subset S$*, there exists a compact set* $L \subset S$ *such that*

1. $\tilde{L} \subset L$*,*
2. $\xi(\partial L^2) = 0$ *almost surely, i.e.* $\mathbb{P}(\xi(\partial L^2) > 0) = 0$*.*

Here, ∂X represents the boundary of a set X, and it should be noted that L^2 is also compact in S^2. The motivation here is that boundary issues break the continuity conditions that are used in the proof, and the general compact

sets from the radial tightness condition do not in general have good boundary behavior, so we may need to enlarge them via this lemma.

We now give the first proposition that we need for the weak convergence theorem. It guarantees that if some functional on regular adjacency measures depends only on the restrictions to such a set L^2, then the functional is continuous at all adjacency measures that have no mass on the boundary of L^2.

▶ **Proposition 4.4** Localizable Functionals. *Consider a square LC Polish space S^2. Let $F : \mathcal{L} \subseteq \mathcal{E}^\circ(S^2) \to \mathbb{R}$ be a functional. Suppose that there exists some compact set $L \subset S$, such that for all*

$$(\tau, x), (\tau', x') \in \mathcal{L}$$

with $x, x' \in L$ and

$$G(\tau|_{L^2}, x) \cong G(\tau'|_{L^2}, x'),$$

it holds that $F(\tau, x) = F(\tau', x')$. Then F is continuous at all $(\tau, x) \in \mathcal{L}$ with $x \in L$ and $\tau(\partial L^2) = 0$. Furthermore, the functional

$$(\tau, x) \mapsto \mathbf{1}_{\{x \in L\}} F(\tau, x)$$

is continuous at all $(\tau, x) \in \mathcal{L}$ with $\tau(\partial L^2) = 0$.

The restriction notation $\tau|_X$ is

$$\tau|_X := \sum_{(x,y) \in \tau \cap X} \delta_{(x,y)}. \tag{1}$$

We will use this result to conclude continuity of specially chosen indicator functionals in the upcoming proofs. Lastly, we give an alternative characterization of spatial tightness in the context where the sequence is known to converge, which is also critical to the proof of the theorem.

▶ **Proposition 4.5** Alternative Characterization of Radial Tightness. *Consider a sequence of regular rooted adjacency processes $(\xi_n, p_n)_{n \in \mathbb{N}} \in \mathcal{E}^\circ(S^2)$ on a square LC Polish space S^2. Suppose that there exists some $(\xi_\infty, p_\infty) \in \mathcal{E}^\circ(S^2)$ for which*

$$(\xi_n, p_n) \Rightarrow (\xi_\infty, p_\infty)$$

in $\mathcal{E}^\circ(S^2)$. Then, $(\xi_n, p_n)_{n \in \mathbb{N}}$ is radially tight if and only if for every $r \in \mathbb{N}$, it holds that

$$|V_r(\xi_n, p_n)| \Rightarrow |V_r(\xi_\infty, p_\infty)|$$

as random counts in $\mathbb{N}$.

The intuition behind this is sensible, as we know that all of these counts will be almost surely finite by the assumed process convergence, and thus convergence of counts prevents vertex escape. Of course, there are many radially tight sequences that do not converge, and thus do not have this property. However, we will only use radial tightness in contexts where process convergence is assumed or also shown, so this condition is immensely useful for us. We give one final small lemma that makes the upcoming arguments cleaner.

▶ **Lemma 4.6.** *Consider a sequence $(\xi_n, p_n)_{n\in\mathbb{N}} \in \mathcal{E}^\circ(S^2)$ on a square LC Polish space S^2. Suppose that there exists some $(\xi_\infty, p_\infty) \in \mathcal{E}^\circ(S^2)$ for which*

$$G(\xi_n, p_n) \Rightarrow G(\xi_\infty, p_\infty).$$

Then,

$$|V_r(\xi_n, p_n)| \Rightarrow |V_r(\xi_\infty, p_\infty)|$$

as random counts in $\mathbb{N}$.

With all of the above in mind, we can give the proof of Theorem 4.1, contingent on the aforementioned results.

Proof of the Weak Convergence Theorem (Theorem 4.1)

Proof. ($\Leftarrow$) Assume radial tightness. By [18, Def. 1.3], weak convergence in $\mathcal{G}_*$ is equivalent to the condition that for every finite rooted graph $\tilde{G} \in \mathcal{G}_*$ and every $r \in \mathbb{N}$, it holds that

$$\mathbb{P}(B_r(\xi_n, p_n) \cong \tilde{G}) \longrightarrow \mathbb{P}(B_r(\xi_\infty, p_\infty) \cong \tilde{G}). \tag{2}$$

Fix $r \in \mathbb{N}$ and $\tilde{G} \in \mathcal{G}_*$. Then, define the functional $\Phi : \mathcal{E}^\circ(S^2) \to \{0, 1\}$ as

$$\Phi(\tau, x) = \mathbf{1}_{\{B_r(\tau,x)\cong\tilde{G}\}}.$$

Equation (2) follows directly from

$$\mathbb{E}(\Phi(\xi_n, p_n)) \to \mathbb{E}(\Phi(\xi_\infty, p_\infty)),$$

so it suffices to show that this convergence of expectations holds. The proof will proceed via the following steps:

1. We choose a compact set $L \subseteq S$ so that it satisfies radial tightness at radius r and so that $\xi_\infty(\partial L^2) = 0$ almost surely.
2. We replace the rooted-ball indicator Φ with a localized indicator Ψ obtained by restricting the adjacency measure to L^2.
3. By Proposition 4.4, (ξ_∞, p_∞) is almost surely a continuity point of the localized indicator Ψ. Since the inputs converge weakly, Proposition 2.1 then gives convergence of the expectations of the localized indicators.
4. Radial tightness gives that, with arbitrarily high probability, the rooted r-ball has all its vertices in L. On this event, the restricted and unrestricted indicators agree, so the error from replacing Φ by Ψ is arbitrarily small.
5. Therefore, convergence of the expectations of Ψ transfers to convergence of the expectations of Φ.

We proceed formally by first constructing L. Fix $\varepsilon > 0$. By radial tightness, there exist a compact set $\tilde{L} \subset S$ and an index $N \in \mathbb{N}$ such that for all $n \geq N$,

$$\mathbb{P}(V_r(\xi_n, p_n) \subset \tilde{L}) \geq 1 - \varepsilon.$$

Applying Lemma 4.3 to the compact set $\tilde{L} \subset S$ and the limiting process ξ_∞ gives a compact set $L \supset \tilde{L}$ with $\xi_\infty(\partial L^2) = 0$ almost surely. Since $\tilde{L} \subset L$, for all $n \geq N$ we have

$$\mathbb{P}(V_r(\xi_n, p_n) \subset L) \geq \mathbb{P}(V_r(\xi_n, p_n) \subset \tilde{L}) \geq 1 - \varepsilon$$

as well. Thus, the enlarged set L still satisfies the required radial tightness bound, while L^2 also has ξ_∞-null boundary. Define the events

$$A_n := \{V_r(\xi_n, p_n) \subset L\}, \qquad A_\infty := \{V_r(\xi_\infty, p_\infty) \subset L\}.$$

Then $\mathbb{P}(A_n) \geq 1 - \varepsilon$ for all $n \geq N$, and $\mathbb{P}(A_\infty) \geq 1 - \varepsilon$ by Lemma 4.2. Define the truncated indicator functional $\Psi : \mathcal{E}^\circ(S^2) \to \{0, 1\}$ as

$$\Psi(\tau, x) = \mathbf{1}_{x \in L} \Phi(\tau|_{L^2}, x).$$

Evidently the codomain here is finite, and thus Ψ is bounded. Furthermore, this functional is localizable, since isomorphic graphs also have the same isomorphism relation to the chosen test graph $\tilde{G}$. Hence, by Proposition 4.4, Ψ is continuous at every (τ, x) satisfying $\tau(\partial L^2) = 0$. Since we chose L so that $\xi_\infty(\partial L^2) = 0$ almost surely, it follows that (ξ_∞, p_∞) is almost surely a continuity point of Ψ. By Proposition 2.1 we can then identify that

$$\mathbb{E}(\Psi(\xi_n, p_n)) \to \mathbb{E}(\Psi(\xi_\infty, p_\infty)). \tag{3}$$

Note then that

$$\Psi(\tau, x) \neq \Phi(\tau, x) \implies V_r(\tau, x) \not\subset L$$

by simple contradiction. Indeed, if we had $V_r(\tau, x) \subset L$, then every vertex and edge needed to determine the rooted r-ball around x would be contained in the restriction $\tau|_{L^2}$. Hence, the restricted and unrestricted constructions would give the same rooted r-ball, and it would hold that $\Psi(\tau, x) = \Phi(\tau, x)$. Therefore, for all $n \geq N$,

$$\mathbb{P}(\Psi(\xi_n, p_n) \neq \Phi(\xi_n, p_n)) \leq \mathbb{P}(V_r(\xi_n, p_n) \not\subset L) < \varepsilon$$

and consequently, via [22, Thm. 5.3(v)], it holds that for all $n \geq N$,

$$|\mathbb{E}(\Phi(\xi_n, p_n)) - \mathbb{E}(\Psi(\xi_n, p_n))| < \varepsilon.$$

Moreover, since we also have that $\mathbb{P}(A_\infty) \geq 1 - \varepsilon$, it holds that

$$|\mathbb{E}(\Phi(\xi_\infty, p_\infty)) - \mathbb{E}(\Psi(\xi_\infty, p_\infty))| < \varepsilon$$

as well. Thus, for all $n \geq N$ we have via the triangle inequality that

$$\begin{aligned} |\mathbb{E}(\Phi(\xi_n, p_n)) - \mathbb{E}(\Phi(\xi_\infty, p_\infty))| < {} & |\mathbb{E}(\Phi(\xi_n, p_n)) - \mathbb{E}(\Psi(\xi_n, p_n))| \\ & + |\mathbb{E}(\Psi(\xi_n, p_n)) - \mathbb{E}(\Psi(\xi_\infty, p_\infty))| \\ & + |\mathbb{E}(\Psi(\xi_\infty, p_\infty)) - \mathbb{E}(\Phi(\xi_\infty, p_\infty))| \\ < {} & 2\varepsilon + |\mathbb{E}(\Psi(\xi_n, p_n)) - \mathbb{E}(\Psi(\xi_\infty, p_\infty))|. \end{aligned}$$

The remainder term vanishes in the limit by Equation (3), and thus since ε was arbitrary, it holds that

$$\mathbb{E}(\Phi(\xi_n, p_n)) \to \mathbb{E}(\Phi(\xi_\infty, p_\infty)),$$

and we conclude the implication due to r and $\tilde{G}$ being arbitrary. Note that this first implication is independent of Proposition 4.5, and will in fact be used to prove said proposition.

($\Rightarrow$) Assume $G(\xi_n, p_n) \Rightarrow G(\xi_\infty, p_\infty)$. Then, Lemma 4.6 followed by Proposition 4.5 conclude the proof. ◀

4.2 Propositions

Below we prove the main technical inputs used above, starting with Proposition 4.4. The proof relies on a structural feature of adjacency measures that originally motivated the inclusion of the diagonal in Section 3.1: small spatial changes of an adjacency measure should not change the graph it constructs, provided that neighboring atoms can be matched.

To make this precise, we replace the informal language of perturbations by a count-based comparison. When two rooted adjacency measures are sufficiently close in the right local sense, their atoms can be matched so that adjacency and rooting are preserved. The corresponding rooted graphs should then be isomorphic.

The mechanism used to enforce this comparison is what we informally call a *counting family*: a countable collection of Borel sets, each with a corresponding reference count. We show that, around any regular rooted adjacency measure, one can choose such a family so that every other regular rooted adjacency measure matching all reference counts, and with a matching root, corresponds to the same rooted isomorphism class.

▶ **Lemma 4.7** Local Invariance of the Graph Construction. *Let (τ, x) be a regular adjacency measure on a square LC Polish space S^2. Then there exists a countable collection of Borel sets $(B_k)_{k\in\mathbb{N}}$ with τ-empty boundary and reference counts $(b_k)_{k\in\mathbb{N}} \in \{0,1\}$ such that*

$$\tau(B_k) = b_k \text{ for all } k \in \mathbb{N}$$

and such that for all $(\tau', x') \in \mathcal{E}^\circ(S^2)$ satisfying $\tau'(B_k) = b_k$ for all $k \in \mathbb{N}$ and $\{(x,x),(x',x')\} \subseteq B_{k_0}$ for some $k_0 \in \mathbb{N}$, it holds that

$$G(\tau, x) \cong G(\tau', x').$$

If the adjacency measure is finite, then one may choose the family to only have a finite number of Borel sets.

Contingent on this lemma, the proof of Proposition 4.4 is simple, and is given below. In short, the proof restricts a rooted adjacency measure to a compact region, and then applies Lemma 4.7 to build finitely many boundary-safe tests and a root neighborhood that force the same local rooted graph structure, making the functional locally constant there. Continuity follows because convergence guarantees those finite tests and root placement eventually agree exactly, so any convergent sequence eventually remains in that neighborhood and the functional values converge.

Proof of Proposition 4.4

Proof. Fix a square LC Polish space S^2, a subset $\mathcal{L} \subseteq \mathcal{E}^\circ(S^2)$, and a functional $F : \mathcal{L} \to \mathbb{R}$. Fix a compact set L satisfying the assumptions. Take $(\tau, x) \in \mathcal{L}$ with $\tau(\partial L^2) = 0$ and $x \in L$. Then, we must have

$$(x, x) \notin \partial L^2.$$

Hence, (x, x) is an interior point of L^2. Since τ is boundedly-finite and L is compact, the restricted measure

$$\tau|_{L^2}$$

is a finite adjacency measure on the compact Polish space L^2. Indeed, simplicity and fiber-boundedness are unchanged by restriction, and the diagonal is maintained because for any $(x, y) \in \tau|_{L^2}$, necessarily $x, y \in L$, and thus $(x, x), (y, y) \in L^2$ and the diagonal is preserved. Applying Lemma 4.7 to $(\tau|_{L^2}, x)$, viewed as a regular rooted adjacency measure on L^2, yields a finite (due to finiteness of the adjacency measure and the compact space L^2) family of Borel sets

$$B_1, \ldots, B_m \subseteq L^2$$

and reference counts $b_1, \ldots, b_m \in \{0, 1\}$ such that:

- $\tau|_{L^2}(\partial B_k) = 0$ for all k
- $\tau|_{L^2}(B_k) = b_k$ for all k

and whenever (τ', x') is a regular rooted adjacency measure on L^2 satisfying

$$\tau'(B_k) = b_k \qquad \text{for all } k \in \{1, \ldots, m\}$$

and

$$\{(x, x), (x', x')\} \subseteq B_{k_0}$$

for some $k_0 \in \{1, \ldots, m\}$, then

$$G(\tau|_{L^2}, x) \cong G(\tau', x')$$

through rooted isomorphism. Moreover, since (x, x) is an interior point of L^2, the construction in Lemma 4.7 may be carried out so that the distinguished set B_{k_0} containing (x, x) is open in S^2 and satisfies

$$B_{k_0} \subseteq L^2.$$

Define

$$U_{(\tau,x)} := \Big\{(\tau', x') \in \mathcal{L} \;\Big|\; \tau'|_{L^2}(B_k) = \tau|_{L^2}(B_k) \text{ for all } k \in \{1, \dots, m\}, \\ (x', x') \in B_{k_0}\Big\}.$$

This set contains (τ, x). Now fix $(\tau', x') \in U_{(\tau,x)}$. Since $(x', x') \in B_{k_0} \subseteq L^2$, we have $x' \in L$, and since B_{k_0} was chosen open we have $x' \notin \partial L$. Since the root is an interior point and L is compact, the restricted measure $(\tau'|_{L^2}, x')$ is a finite rooted adjacency measure on L^2 via the earlier reasoning. By construction,

$$\tau'|_{L^2}(B_k) = \tau|_{L^2}(B_k) = b_k \qquad \text{for all } k,$$

and

$$\{(x, x), (x', x')\} \subseteq B_{k_0}.$$

Therefore, Lemma 4.7 applied on the ambient space L^2, gives

$$G(\tau|_{L^2}, x) \cong G(\tau'|_{L^2}, x')$$

and the hypothesis on F, this implies

$$F(\tau', x') = F(\tau, x).$$

Hence, F is constant on $U_{(\tau,x)}$. It remains to show that if $(\tau_n, x_n) \to (\tau, x)$ in $\mathcal{L}$, then eventually

$$(\tau_n, x_n) \in U_{(\tau,x)}.$$

Fix such a sequence. Since $\tau(\partial L^2) = 0$, the restriction map

$$R_{L^2} : \mu \mapsto \mu|_{L^2}$$

is continuous at τ via [2, Prop. 2.8]. Hence,

$$\tau_n|_{L^2} \to \tau|_{L^2}$$

in $\mathcal{E}(L^2)$. Since also $x_n \to x$ in S and $x \notin \partial L$, we obtain

$$(\tau_n|_{L^2}, x_n) \to (\tau|_{L^2}, x)$$

in $\mathcal{E}^\circ(L^2)$. Each B_k is a Borel subset of the compact space L^2, hence relatively compact in L^2, and satisfies

$$\tau|_{L^2}(\partial B_k) = 0.$$

Therefore, vague convergence on L^2 gives

$$\tau_n|_{L^2}(B_k) \to \tau|_{L^2}(B_k) \qquad \text{for each } k \in \{1, \dots, m\}.$$

Since these quantities are integer-valued, for each k there exists some $N_k \in \mathbb{N}$ such that

$$\tau_n|_{L^2}(B_k) = \tau|_{L^2}(B_k)$$

for all $n \geq N_k$. Let

$$M_1 := \max_{1 \leq k \leq m} N_k.$$

Then for all $n \geq M_1$,

$$\tau_n|_{L^2}(B_k) = \tau|_{L^2}(B_k) \qquad \text{for all } k \in \{1, \ldots, m\}.$$

Since B_{k_0} is open in S^2 and contains (x, x), and since

$$(x_n, x_n) \to (x, x)$$

in S^2, there exists an $M_2 \in \mathbb{N}$ such that

$$(x_n, x_n) \in B_{k_0}$$

for all $n \geq M_2$. Therefore for all $n \geq \max(M_1, M_2)$ we have both

$$\tau_n|_{L^2}(B_k) = \tau|_{L^2}(B_k) \qquad \text{for all } k \in \{1, \ldots, m\}$$

and

$$(x_n, x_n) \in B_{k_0}.$$

Hence, $(\tau_n, x_n) \in U_{(\tau,x)}$ for all sufficiently large n. Since F is constant on $U_{(\tau,x)}$, it follows that

$$F(\tau_n, x_n) = F(\tau, x)$$

for all sufficiently large n. Therefore,

$$F(\tau_n, x_n) \to F(\tau, x),$$

and thus F is continuous at (τ, x). For the second claim, define

$$\Psi(\tau, x) := \mathbf{1}_{\{x \in L\}} F(\tau, x).$$

Fix $(\tau, x) \in \mathcal{L}$ with $\tau(\partial L^2) = 0$. If $x \in L$, continuity of Ψ at (τ, x) follows from the first part. If $x \notin L$, then since L is compact and hence closed, $S \setminus L$ is an open neighborhood of x. Thus for any sequence $(\tau_n, x_n) \to (\tau, x)$, we have $x_n \notin L$ for all sufficiently large n, and therefore

$$\Psi(\tau_n, x_n) = 0 = \Psi(\tau, x)$$

eventually. Hence, Ψ is continuous at every $(\tau, x) \in \mathcal{L}$ with $\tau(\partial L^2) = 0$. This concludes the proof. ◀

Of course, Proposition 4.4 is useful in a lot of cases. The general strategy for its application is to restrict a class of functionals that you want a result on to such a compact set, usually chosen to be compatible with radial tightness, and then use the tightness to collapse the error between the truncated functionals and the originals. This was already seen in the proof of Theorem 4.1, but will also be seen in the proof below of our other core proposition, Proposition 4.5, where we use a similar truncation strategy on the vertex count functional.

Proof of Proposition 4.5

Proof. ($\Leftarrow$) First assume the vertex count convergence. We prove radial tightness by induction on the graph radius. We need to show that for every radius $r \in \mathbb{N}$ and every $\varepsilon > 0$, there exists a compact set $L_r \subset S$ and some $N_r \in \mathbb{N}$ for which for all $n \geq N_r$,

$$\mathbb{P}(V_r(\xi_n, p_n) \not\subset L_r) < \varepsilon \tag{4}$$

We use the disambiguations L_r and N_r here to make the induction argument cleaner. In short, we essentially want to show that if Equation (4) holds for every $\varepsilon > 0$ when we fix some $r \in \mathbb{N}$, it also holds for every $\varepsilon > 0$ at $r + 1$.

(Base Case) Fix $r = 0$ and $\varepsilon > 0$. Evidently, $V_0(\xi_n, p_n) = \delta_{p_n}$. Since p_∞ is S-valued, and S is Polish, it holds by [22, Thm. 13.6] that there exists some relatively compact open set $L \subset S$ for which

$$\mathbb{P}(p_\infty \in L) > 1 - \varepsilon$$

Since $p_n \Rightarrow p_\infty$ is assumed, Proposition 2.1 applied to the open set L gives

$$\liminf_{n\to\infty} \mathbb{P}(p_n \in L) \geq \mathbb{P}(p_\infty \in L) > 1 - \varepsilon.$$

Thus, there exists some $N_0 \in \mathbb{N}$ such that for all $n \geq N_0$, letting $L_0 = \overline{L}$, it holds that

$$\mathbb{P}(V_0(\xi_n, p_n) \subset L_0) = \mathbb{P}(p_n \in L_0) \geq 1 - \varepsilon,$$

so since ε was arbitrary Equation (4) holds at $r = 0$ via L_0 and N_0.

(Inductive Case) Fix $r \in \mathbb{N}$. We show that the radial tightness bound at radius r can be made arbitrarily small. Let $\varepsilon > 0$. By the inductive hypothesis, there exist a compact set $L_{r-1} \subset S$ and an index $N_{r-1} \in \mathbb{N}$ such that for all $n \geq N_{r-1}$,

$$\mathbb{P}(V_{r-1}(\xi_n, p_n) \not\subset L_{r-1}) < \varepsilon.$$

By bounded finiteness, $V_r(\xi_\infty, p_\infty)$ is almost surely finite. Hence, there exists $K \in \mathbb{N}$, $K \geq 1$, such that

$$\mathbb{P}(|V_r(\xi_\infty, p_\infty)| > K) < \varepsilon.$$

Fix any $c \in S$, and define

$$R_r := \sup\{d_S(c, v) : v \in V_r(\xi_\infty, p_\infty)\}.$$

Since $V_r(\xi_\infty, p_\infty)$ is almost surely finite, $R_r < \infty$ almost surely. Hence, there exists $M < \infty$ such that

$$\mathbb{P}(R_r > M) < \frac{\varepsilon}{K}.$$

Set

$$L := \overline{B_S(c, M)} \cup L_{r-1}.$$

By properness of d_S, $\overline{B_S(c,M)}$ is compact, so L is compact and $L_{r-1} \subset L$. Moreover,

$$\mathbb{P}(V_r(\xi_\infty, p_\infty) \not\subset L) \leq \mathbb{P}(R_r > M) < \frac{\varepsilon}{K}.$$

Applying Lemma 4.3 to L and ξ_∞, choose a compact set $L_r \supset L$ such that

$$\xi_\infty(\partial L_r^2) = 0 \qquad \text{almost surely.}$$

Since $L \subset L_r$, we still have

$$\mathbb{P}(V_r(\xi_\infty, p_\infty) \not\subset L_r) < \frac{\varepsilon}{K}. \tag{5}$$

For this L_r and r, define the restricted vertex count functional

$$\Phi : \mathcal{E}^\circ(S^2) \to \mathbb{N}, \qquad \Phi(\tau, x) := \mathbf{1}_{\{x \in L_r\}} |V_r(\tau|_{L_r^2}, x)|.$$

This functional is localizable, since rooted graph isomorphism preserves the number of vertices in the radius-r ball. Hence, by Proposition 4.4, Φ is continuous at every (τ, x) satisfying $\tau(\partial L_r^2) = 0$. Since L_r was chosen so that $\xi_\infty(\partial L_r^2) = 0$ almost surely, it follows that

$$\Phi(\xi_n, p_n) \Rightarrow \Phi(\xi_\infty, p_\infty).$$

Define the unrestricted and restricted counts

$$U_n := |V_r(\xi_n, p_n)|, \qquad T_n := \Phi(\xi_n, p_n),$$

and

$$U_\infty := |V_r(\xi_\infty, p_\infty)|, \qquad T_\infty := \Phi(\xi_\infty, p_\infty).$$

By assumption, $U_n \Rightarrow U_\infty$, and by the preceding paragraph, $T_n \Rightarrow T_\infty$. For all $n \in \mathbb{N}$, restriction can only delete vertices and edges, so $T_n \leq U_n$ almost surely. The same holds at the limit. Moreover, on the event

$$E := \{V_r(\xi_\infty, p_\infty) \subset L_r\},$$

the restricted and unrestricted radius-r balls agree, so

$$T_\infty = U_\infty \qquad \text{on } E.$$

Hence

$$\{T_\infty < U_\infty\} \subseteq E^c,$$

and therefore, by Equation (5),

$$\mathbb{P}(T_\infty < U_\infty) < \frac{\varepsilon}{K}.$$

Now let

$$q_K(j) := \min(j, K).$$

Since $\mathbb{N}$ has the discrete topology, both q_K and $\mathbf{1}_{\{j>K\}}$ are bounded continuous functions on $\mathbb{N}$. Thus, by Proposition 2.1,

$$\mathbb{E}[q_K(U_n)] \to \mathbb{E}[q_K(U_\infty)], \qquad \mathbb{E}[q_K(T_n)] \to \mathbb{E}[q_K(T_\infty)],$$

and

$$\mathbb{P}(U_n > K) \to \mathbb{P}(U_\infty > K).$$

Consider the following decomposition:

$$\{T_n < U_n\} \subseteq \{U_n > K\} \cup \{T_n < U_n,\ U_n \le K\}.$$

On $\{T_n < U_n,\ U_n \le K\}$, we have $T_n < U_n \le K$, hence,

$$q_K(U_n) - q_K(T_n) = U_n - T_n \ge 1.$$

Since also $q_K(U_n) - q_K(T_n) \ge 0$ a.s. (because $T_n \le U_n$ a.s. and q_K is non-decreasing),

$$\mathbf{1}_{\{T_n<U_n,\ U_n\le K\}} \le q_K(U_n) - q_K(T_n) \quad \text{a.s.}$$

Therefore,

$$\begin{aligned}
\mathbb{P}(T_n < U_n) &\le \mathbb{P}(U_n > K) + \mathbb{P}(T_n < U_n,\ U_n \le K) \\
&= \mathbb{P}(U_n > K) + \mathbb{E}\big[\mathbf{1}_{\{T_n<U_n,\ U_n\le K\}}\big] \\
&\le \mathbb{P}(U_n > K) + \mathbb{E}[q_K(U_n) - q_K(T_n)]\,.
\end{aligned}$$

Taking lim sup,

$$\begin{aligned}
\limsup_{n\to\infty} \mathbb{P}(T_n < U_n) &\le \limsup_{n\to\infty} \Big(\mathbb{P}(U_n > K) + \mathbb{E}[q_K(U_n) - q_K(T_n)]\Big) \\
&= \mathbb{P}(U_\infty > K) + \mathbb{E}[q_K(U_\infty) - q_K(T_\infty)] \qquad \text{(Portm.)} \\
&\le \mathbb{P}(U_\infty > K) + \mathbb{E}\big[K\,\mathbf{1}_{\{T_\infty<U_\infty\}}\big] \\
&= \mathbb{P}(U_\infty > K) + K\,\mathbb{P}(T_\infty < U_\infty) \\
&< \varepsilon + \varepsilon = 2\varepsilon.
\end{aligned}$$

Now observe that if

$$V_{r-1}(\xi_n, p_n) \subset L_{r-1} \qquad \text{but} \qquad V_r(\xi_n, p_n) \not\subset L_r,$$

then some new radius-r vertex lies outside L_r. Restriction to L_r^2 removes this vertex and cannot add vertices, so the restricted radius-r count is strictly smaller. Hence, denoting $V_{n,r-1} := V_{r-1}(\xi_n, p_n)$ and $V_{n,r} := V_r(\xi_n, p_n)$,

$$\begin{aligned}
\{V_{n,r-1} \subset L_{r-1}\} \cap \{V_{n,r} \not\subset L_r\} &\subseteq \{T_n < U_n\}, \\
\{V_{n,r} \not\subset L_r\} &\subseteq \{T_n < U_n\} \cup \{V_{n,r-1} \not\subset L_{r-1}\}.
\end{aligned}$$

Therefore,

$$\begin{aligned}
\limsup_{n\to\infty} \mathbb{P}(V_{n,r} \not\subset L_r) &\le \limsup_{n\to\infty} \mathbb{P}(T_n < U_n) + \limsup_{n\to\infty} \mathbb{P}(V_{n,r-1} \not\subset L_{r-1}) \\
&\le 2\varepsilon + \varepsilon = 3\varepsilon.
\end{aligned}$$

Since $\varepsilon > 0$ was arbitrary, the radial tightness bound at radius r follows. By induction, we conclude that $(\xi_n, p_n)_{n\in\mathbb{N}}$ is a radially tight sequence.

($\Rightarrow$) Conversely, if we assume radial tightness, we can use the backward implication ($\Leftarrow$) of Theorem 4.1, which is independent of this proposition, to conclude convergence of the corresponding graphs, and then applying Lemma 4.6 gives the desired condition. ◀

4.3 Lemmas

We now give the proof of the relevant lemmas. We begin with the proof of Lemma 4.7 as it is perhaps the most fundamental, and the argument is worth illustrating. Essentially, we construct a sufficiently small open ball around each vertex. These function as separating balls that isolate the vertices. Each of these has a reference count of one. Then, we take the complement of the union of these balls in the ambient space, which is not open but is Borel, and assign it a reference count of zero. This essentially mimics the idea of "place a vertex close to each original vertex, and have no other vertices floating around anywhere". The proof separates the atoms of an adjacency measure τ by small disjoint test neighborhoods and adds the complement as a zero-count set. Any τ' with the same local counts must therefore have exactly one corresponding atom in each neighborhood and no extra atoms. Matching diagonal atoms defines a bijection of vertices, and the chosen neighborhoods force edge atoms to match their vertices under this bijection. Hence, the induced graph isomorphism preserves the root, and $G(\tau, x) \cong G(\tau', x')$.

Proof of Lemma 4.7

Proof. Since τ is a simple boundedly-finite counting measure on S^2, every atom of τ is isolated, and the set of atoms is countable. Thus, for each $z \in \tau$ we may choose an open neighborhood $A_z \subseteq S^2$ of z such that:

- $\tau(A_z) = 1$;
- $\tau(\partial A_z) = 0$;
- the family $(A_z)_{z\in\tau}$ is pairwise disjoint.

If τ is finite, this is of course only a finite family. By shrinking if necessary, we may additionally ensure that for every $(a, b) \in \tau$,

$$(u, v) \in A_{(a,b)} \implies (u, u) \in A_{(a,a)} \text{ and } (v, v) \in A_{(b,b)}.$$

Indeed, this follows from continuity of the maps

$$(u, v) \mapsto (u, u) \qquad \text{and} \qquad (u, v) \mapsto (v, v)$$

on S^2, together with the fact that $A_{(a,a)}$ and $A_{(b,b)}$ are open neighborhoods of (a, a) and (b, b) respectively. Now define

$$D^* := S^2 \setminus \bigcup_{z\in\tau} A_z.$$

Then D^* is a Borel set. Also, since every atom of τ already lies in some A_z, we have that $\tau(D^*) = 0$. Moreover, every atom of τ lies in the interior of $\bigcup_{z\in\tau} A_z$, and hence no atom of τ lies on the boundary of this union, and thus $\tau(\partial D^*) = 0$ as well.

We now take as our family of Borel sets all sets A_z and the single set D^*. Re-index these arbitrarily as $(B_k)_{k\in K}$ where K is countable in general and finite when τ is finite. Define

$$b_k := \begin{cases} 1, & \text{if } B_k = A_z \text{ for some } z \in \tau, \\ 0, & \text{if } B_k = D^*. \end{cases}$$

Then each B_k is Borel, satisfies $\tau(B_k) \leq 1$, has τ-null boundary, and also satisfies $\tau(B_k) = b_k$. Now let (τ', x') be a regular rooted adjacency measure satisfying

$$\tau'(B_k) = b_k \qquad \text{for all } k \in K$$

and

$$\{(x, x), (x', x')\} \subseteq B_k \qquad \text{for some } k \in K.$$

Since $\tau'(D^*) = 0$, every atom of τ' lies in

$$\bigcup_{z\in\tau} A_z.$$

Since the sets A_z are pairwise disjoint and each satisfies $\tau'(A_z) = 1$, it follows that each A_z contains exactly one atom of τ', and every atom of τ' lies in exactly one such set. We first define the vertex correspondence. For each $a \in V(\tau)$, let

$$(\phi(a), \phi(a))$$

be the unique atom of τ' in $A_{(a,a)}$. This is well-defined because $(a, a) \in \tau$ and $\tau'(A_{(a,a)}) = 1$. Thus, we obtain a map $\phi : V(\tau) \to V(\tau')$ which is injective because the sets $A_{(a,a)}$ are pairwise disjoint.

To see surjectivity, let $b \in V(\tau')$. Then $(b, b) \in \tau'$, so by $\tau'(D^*) = 0$ the point (b, b) lies in some unique set $A_{(u,v)}$. By the defining property of the sets A_z,

$$(b, b) \in A_{(u,u)} \cap A_{(v,v)},$$

but since the sets $A_{(a,a)}$ are pairwise disjoint, we must have $u = v$. Hence, $(b, b) \in A_{(u,u)}$ and therefore $b = \phi(u)$. Thus, ϕ is bijective.

We now show that ϕ preserves edges. Let $(a, b) \in \tau$. Since $\tau'(A_{(a,b)}) = 1$, there is a unique atom $(u, v) \in \tau' \cap A_{(a,b)}$. By the defining property of $A_{(a,b)}$,

$$(u, u) \in A_{(a,a)} \qquad \text{and} \qquad (v, v) \in A_{(b,b)}.$$

By definition of ϕ, this implies that $u = \phi(a)$ and $v = \phi(b)$. Hence, $(\phi(a), \phi(b)) \in \tau'$.

Conversely, let $(u, v) \in \tau'$. Since $\tau'(D^*) = 0$, the point (u, v) lies in some unique set $A_{(a,b)}$, and by the same argument, $u = \phi(a)$ and $v = \phi(b)$. Since $A_{(a,b)}$ contains exactly one atom of τ', namely (u, v), and since $\tau(A_{(a,b)}) = 1$ with $(a, b) \in A_{(a,b)}$, we conclude that $(a, b) \in \tau$. Therefore,

$$(a, b) \in \tau \iff (\phi(a), \phi(b)) \in \tau'.$$

So ϕ induces a graph isomorphism $G(\tau) \cong G(\tau')$. Finally, since

$$\{(x, x), (x', x')\} \subseteq B_k$$

for some $k \in K$, and the only set in our family with truth value 0 is D^*, this set must have truth value 1. Hence, it is of the form $B_k = A_z$ for some $z \in \tau$. Since $(x, x) \in A_z$ and the diagonal sets $A_{(a,a)}$ are pairwise disjoint, we must have $z = (x, x)$. Thus, $(x', x') \in A_{(x,x)}$ and by definition of ϕ this implies that

$$\phi(x) = x'$$

Hence the induced graph isomorphism is rooted, and we conclude that $G(\tau, x) \cong G(\tau', x')$. ◀

We now give the proof of the remaining, simpler lemmas in order of occurrence. First Lemma 4.2, then Lemma 4.3, and then lastly Lemma 4.6. All of these proofs are straightforward and technical, and thus we don't give a detailed overview of the proof strategies.

Proof of Lemma 4.2

Proof. Define

$$C_{r,L} := \{(\tau, p) \in \mathcal{E}^\circ(S^2) : V_r(\tau, p) \subset L\}.$$

We first show that $C_{r,L}$ is closed. Let $(\tau_n, p_n) \in C_{r,L}$ and $(\tau_n, p_n) \to (\tau, p)$ in $\mathcal{E}^\circ(S^2)$. Assume for contradiction that $(\tau, p) \notin C_{r,L}$. Then there exists

$$x \in V_r(\tau, p) \setminus L.$$

Since L is compact (hence closed), choose an open neighborhood $O_x \subset S$ with $x \in O_x$ and $O_x \cap L = \varnothing$. Now form the finite rooted adjacency measure obtained by restricting τ to the radius-r vertex cloud:

$$\tau^{(r)} := \tau|_{\{(u,v) \in S^2 : u \in V_r(\tau,p),\ v \in V_r(\tau,p)\}}.$$

Reusing the separating-neighborhood construction from the proof of Lemma 4.7, choose disjoint open sets $(A_z)_{z \in \tau^{(r)}}$ with $\tau(\partial A_z) = 0$, $\tau(A_z) = 1$, and in particular with $A_{(x,x)} \subseteq O_x$. Because $(\tau_n, p_n) \to (\tau, p)$, the same boundary-null count-stability argument used in the proof of Proposition 4.4 gives, for all sufficiently large n:

$$\tau_n(A_z) = \mathbf{1}_{z \in \tau^{(r)}}, \qquad (p_n, p_n) \in A_{(p,p)}.$$

Hence for each $z \in \tau^{(r)}$, there is a unique corresponding atom of τ_n in A_z. In particular, due to the stability of the vertices and edges via the Lemma 4.7 argumentation, the vertex x yields some sequence $x_n \in V_r(\tau_n, p_n)$ with

$$(x_n, x_n) \in A_{(x,x)}$$

eventually. Therefore, x_n is eventually contained in O_x which is entirely outside L, contradicting $V_r(\tau_n, p_n) \subset L$. So $(\tau, p) \in C_{r,L}$, and $C_{r,L}$ is closed. Now set

$$A_n := \{V_r(\xi_n, p_n) \subset L\} = \{(\xi_n, p_n) \in C_{r,L}\}.$$

By assumption, $\mathbb{P}(A_n) \geq 1 - \varepsilon$ for all $n \geq N$, so

$$\limsup_{n\to\infty} \mathbb{P}(A_n) \geq 1 - \varepsilon.$$

Since $(\xi_n, p_n) \Rightarrow (\xi_\infty, p_\infty)$ and $C_{r,L}$ is closed, Proposition 2.1 gives

$$\mathbb{P}((\xi_\infty, p_\infty) \in C_{r,L}) \geq \limsup_{n\to\infty} \mathbb{P}((\xi_n, p_n) \in C_{r,L}) \geq 1 - \varepsilon.$$

Equivalently,

$$\mathbb{P}(V_r(\xi_\infty, p_\infty) \subset L) \geq 1 - \varepsilon.$$

This proves the claim. ◀

Proof of Lemma 4.3

Proof. Recall that $\tilde{L}$ is compact, and thus we can choose $c \in S$, and identify the real number

$$R_0 := \sup_{y\in\tilde{L}} d_S(c, y) < \infty.$$

Then $\tilde{L} \subseteq \overline{B_S(c, R_0)}$. For $r \geq R_0$, define $L_r := \overline{B_S(c, r)}$ and bad event

$$E_r := \{\xi(\partial L_r^2) > 0\}.$$

We consider some probability space for ξ, denoted $(\Omega, \mathcal{F}, \mathbb{P})$. For each $\omega \in \Omega$, define bad-radius set

$$\mathcal{B}(\omega) := \{r \geq R_0 : \xi(\omega)(\partial L_r^2) > 0\}.$$

Because each realization $\xi(\omega)$ is a simple boundedly finite counting measure on σ-compact S^2, its atom set is countable. If $r \in \mathcal{B}(\omega)$, some atom (u, v) lies in ∂L_r^2, and thus either $d_S(c, u) = r$ or $d_S(c, v) = r$. So $\mathcal{B}(\omega)$ is contained in a countable set of coordinate distances, hence $\mathcal{B}(\omega)$ is countable. Every countable subset of $\mathbb{R}$ is Lebesgue-null, so

$$\lambda(\mathcal{B}(\omega) \cap [R_0, R_0 + 1]) = 0 \quad \text{for all } \omega \in \Omega.$$

Hence,

$$\int_{R_0}^{R_0+1} \mathbb{P}(E_r)\,dr = \mathbb{E}\left[\int_{R_0}^{R_0+1} \mathbf{1}_{\{r\in\mathcal{B}(\omega)\}}\,dr\right] = \mathbb{E}[\lambda(\mathcal{B}(\omega)\cap[R_0,R_0+1])] = 0.$$

This gives that $\mathbb{P}(E_r)=0$ for Lebesgue-a.e. $r\in[R_0,R_0+1]$. Choose such an $r^\star$, and define $L := L_{r^\star}$. Since $r^\star \geq R_0$, we have $\tilde{L}\subseteq L$. Also $\mathbb{P}(\xi(\partial L^2)>0)=0$, i.e.

$$\xi(\partial L^2) = 0 \text{ almost surely.}$$

This proves the claim. ◀

Proof of Lemma 4.6

Proof. Since $G(\xi_n,p_n)\Rightarrow G(\xi_\infty,p_\infty)$ in $\mathcal{G}_*$, Proposition 2.2 may be applied to any continuous functional on $\mathcal{G}_*$. Fix $r\in\mathbb{N}$, and define

$$F_r([G,o]) := |V(B_r(G,o))|.$$

We claim that F_r is continuous. Indeed, if $[G_k,o_k]\to[G,o]$ in the Benjamini–Schramm topology, then by definition of this topology the rooted r-balls $B_r(G_k,o_k)$ and $B_r(G,o)$ are eventually rooted-isomorphic. Hence they eventually have the same number of vertices, so

$$F_r([G_k,o_k]) = F_r([G,o])$$

for all sufficiently large k. Therefore F_r is continuous.

By construction of $G(\xi,p)$, the vertex set of $B_r(G(\xi,p))$ is exactly the spatial vertex cloud $V_r(\xi,p)$. Thus

$$F_r(G(\xi,p)) = |V_r(\xi,p)|.$$

Consequently, Proposition 2.2 gives

$$|V_r(\xi_n,p_n)| = F_r(G(\xi_n,p_n)) \Rightarrow F_r(G(\xi_\infty,p_\infty)) = |V_r(\xi_\infty,p_\infty)|.$$

◀

5 Local Weak Convergence with Uniform Rooting

With general weak convergence established, the question of local weak convergence follows naturally. Here, we take a sequence of finite adjacency processes $(\xi_n)_{n\in\mathbb{N}}$ on a square LC Polish space S^2, and use the uniform rooting kernel $\mathbb{U}$ from Section 3.3 to select a uniform root. Local weak convergence to a limiting random graph G_∞ taking values in $\mathcal{G}_*$ is then defined as the weak convergence

$$G(\xi_n, U_n) \Rightarrow G_\infty$$

where $U_n \sim \mathbb{U}(\xi_n, \cdot)$. Of course, in order to apply Theorem 4.1, we would like to find some limit $(\xi_\infty, p_\infty) \in \mathcal{E}^\circ(S^2)$ giving $G_\infty = G(\xi_\infty, p_\infty)$ for which

$$(\xi_n, U_n) \Rightarrow (\xi_\infty, p_\infty)$$

in $\mathcal{E}^\circ(S^2)$. Unfortunately, the existence of such a limiting object does not hold in general in this case, as uniformly rooting over adjacency measures generally does not give a limiting root distribution as the adjacency measures get infinitely large.

This ties into the larger idea that there are many sequences of rooted regular adjacency processes that have no vague limits, but from which convergent sequences of locally finite rooted graphs emerge. In these cases, the weak convergence theorem cannot be applied immediately.

Even if the graph convergence can be shown via other means, we still generally want to obtain a regular rooted adjacency process that constructs the limit graph, which often does not come cleanly from a local weak graph limit itself. The natural remedy to this is to transform our sequence of adjacency processes into another sequence of adjacency processes that both

1. induces the same sequence of graphs, and
2. has a limit in $\mathcal{E}^\circ(S^2)$.

More formally, we can consider a sequence of isomorphism-preserving transformations $(F_n)_{n\in\mathbb{N}} : \mathcal{E}^\circ(S^2) \to \mathcal{E}^\circ(S^2)$, satisfying for every $n \in \mathbb{N}$ the identity

$$G(F_n(\xi_n, p_n)) \cong G(\xi_n, p_n),$$

and such that $F_n(\xi_n, p_n)$ is a radially tight sequence which converges to some limit $(\xi_\infty, p_\infty) \in \mathcal{E}^\circ(S^2)$. Then the limit would transfer cleanly, giving

$$G(\xi_n, p_n) \Rightarrow G(\xi_\infty, p_\infty).$$

In general, there is no established way to construct such a suitable sequence of transformations. However, in the specific case of local weak convergence, it is possible to leverage deeply established theory. Essentially, local weak convergence asks what a local graph region looks like around a "typical" root. Analogously in the adjacency measure representation, we may want to ask what the local spatial region looks like around a "typical" point.

Fortunately, the latter question already has an answer in the form of Palm theory. It is known that for a point process on a space X, one can derive its *Palm distribution* at a point $x \in X$, and under sufficient regularity conditions on the point process, the point x in this distribution represents a "typical point" of the original distribution. See [23] for a general reference.

It also happens to hold that this Palm conditioning (or more precisely, a slight variant of it) is an isomorphism-preserving transformation when applied to regular adjacency measures. Thus, we can use Palm distributions to transform our uniformly rooted sequences, and then apply Theorem 4.1 to these distributions to obtain a usable local weak limit.

In this section, we develop this perspective, allowing local weak convergence to be studied through rooted adjacency processes. We introduce the Palm transformation and the slight variant of it that is better suited to uniform rooting, explain the finite approximation setting in which these ideas are used, and then arrive at a specialized result for local weak convergence that builds off of Theorem 4.1.

5.1 Palm Transformations on Adjacency Processes

We first define stationarity for point processes. Intuitively, this means that when you "move around" the point process in space, its distribution looks the same. This should not be confused with the stronger condition of joint exchangeability [36, Def. 4.5] in the existing literature, which is a much more restrictive type of spatial symmetry, and is not required for our results.

The full theory requires the consideration of group actions on our LC Polish spaces, and this will be given shortly, but first we give an example for a point process in $\mathbb{R}$. In this case, we say that a point process η is stationary if for every $x \in \mathbb{R}$,

$$\theta_x \eta =^d \eta.$$

Here, the shift operator θ_x applied to a counting measure μ is the map

$$\mu \mapsto \sum_{z \in \mu} \delta_{z+x},$$

which is measurable by [23, Eq. 2.4] and thus can be applied to point processes. In this example, the group action is translation by an element of $\mathbb{R}$, defined via the standard sum. In general, we can do this with any group H and corresponding group action, so long as a few conditions are met. These assumptions are standard in Palm theory on spaces with group actions. We state them explicitly to make clear which structure is being used, but their role should be understood mainly as ensuring that shifts, stationarity, and Palm distributions are well-behaved.

Firstly, H must be a locally compact, second countable topological group. In essence, this means that H is both a group and a locally compact topological

space, where the topology is compatible with the group operations. At a high level, it ensures that there are natural reference measures on H, that random measures can be studied locally on compact sets, and that the relevant conditioning and measurability arguments behave well. These are the basic ingredients needed to define stationarity and Palm distributions without introducing unnecessary technical pathologies.

Secondly, we also would like the group H to be unimodular. This means that its aforementioned reference measures, known as Haar measures, see [23, Eq. 2.1], agree up to normalization. In short, unimodularity prevents extra correction factors from appearing when shifting configurations by group elements. This makes stationarity and Palm identities cleaner and more symmetric.

Lastly, when we have such a group H acting on an LC Polish space S, we want the group action to be measurable and proper. This means that the map

$$H \times S \to S, \qquad (h, s) \mapsto h \cdot s$$

is measurable, and that the action does not move compact sets into each other in an uncontrolled way. More precisely, properness means that for compact sets $K, L \subseteq S$, the set of group elements sending some point of K into L,

$$\{h \in H : hK \cap L \neq \varnothing\},$$

is compact in the group topology. Measurability allows the action to be used with random objects, while properness ensures that local parts of the space are only affected by locally bounded parts of the group. This prevents pathological behavior when defining shifted configurations, stationarity, and Palm-type constructions.

We bundle such a group H and space S together to give us the following definition.

Definition 16 Grouped LC Polish space

Let S be an LC Polish space, and let H be a unimodular, locally compact, second countable topological group. We call (S, H) a *grouped LC Polish space* iff H acts on S measurably and properly. We also attach to H an associated Haar measure λ_H, for which we denote the pushforward to S as $\lambda_{(S,H)}$.

For a point process η on such a space, stationarity can then be cleanly defined as the property that for all $h \in H$,

$$\theta_h \eta =^d \eta, \tag{6}$$

where the shift operator θ_h is the map

$$\mu \mapsto \sum_{z \in \mu} \delta_{h \cdot z}. \tag{7}$$

This is evidently an abstraction from the earlier case in $\mathbb{R}$. However, it should now be evident that our adjacency processes are usually not stationary in the intuitive sense. For example, in the case of translations on $\mathbb{R}^2$, a point process with labels in this space would not be stationary under this action, since diagonal points might not remain on the diagonal. In order to repair this, we work only with pointwise group actions on our square LC Polish spaces. Specifically,

Definition 17 Square Grouped LC Polish space

Let (S, H) be a grouped LC Polish space. The associated *square grouped LC Polish space* is $(S, H)^2$, meaning the space S^2 (taking a proper product metric from S) equipped with the pointwise action

$$h(x, y) := (hx, hy).$$

The measure on such a space is given by

$$\lambda_{(S,H)^2}(A \times B) := \lambda_{(S,H)}(A)\lambda_{(S,H)}(B)$$

for all Borel sets $A, B \subset S$.

Revisiting our counterexample with the pointwise group action of $\mathbb{R}$-translations on $\mathbb{R}^2$ (instead of the action of $\mathbb{R}^2$-translations) indeed preserves the diagonal, and thus it is possible for our adjacency processes to be stationary in this space. It should be noted that square grouped LC Polish spaces are themselves grouped LC Polish spaces.

We also refer to some more examples of these types of group actions in [35, Sec. 17]. With the above established, we can finally give the definition of the Palm distribution, and then show how it benefits from this notion of stationarity. This is given in the form of a Markov kernel called the Palm kernel, which gives for a point in the ambient space a point process law (distribution of counting measures) on that space.

Definition 18 Palm Kernel

Let S be an LC Polish space, and let η be a simple boundedly finite point process on S with intensity measure Λ_η. A *Palm kernel* of η is a Markov kernel

$$\mathbb{P} : S \times \mathcal{B}(\mathcal{N}(S)) \to [0, 1],$$

written as

$$\mathbb{P}^x(A) := \mathbb{P}(x, A),$$

such that for every measurable $f : S \times \mathcal{N}(S) \to [0, \infty)$,

$$\mathbb{E}\left[\sum_{y \in \eta} f(y, \eta)\right] = \int_S \mathbb{E}^x[f(x, \eta)]\, \Lambda_\eta(dx), \tag{8}$$

where

$$\mathbb{E}^x[f(x,\eta)] := \int_{\mathcal{N}(S)} f(x,\zeta)\,\mathbb{P}^x(d\zeta).$$

Since η is boundedly finite in our setup, Λ_η is σ-finite, and the Palm kernel exists and is unique Λ_η-almost everywhere; see [13, p. 462].

Whenever η^x is a point process with law $\mathbb{P}^x$, we write

$$\eta^x \sim \mathbb{P}^x(\cdot)$$

and refer to η^x as a Palm version of η at x. Thus, at a high level, one can also view this as a transformation mapping $(\eta, x) \mapsto \eta^x$. This viewpoint is what we refer to as the *Palm transformation* of η henceforth.

The intensity measure Λ_η as used above is the unique measure on S given by

$$\Lambda_\eta(A) := \mathbb{E}[\eta(A)] \tag{9}$$

which essentially describes how many points the point process is expected to place in each region of space. What the Palm kernel does, in short, is describe the law of the point process as seen from one of its own points. More precisely, $\mathbb{P}^x$ should be thought of as the conditional distribution of η, given that there is a point at x. Since the event "there is a point exactly at x" often has probability zero, this conditional distribution cannot usually be defined by ordinary conditioning. Instead, it is characterized by the identity above.

The left-hand side of Equation (8) averages $f(y,\eta)$ over all points y of the process, while the right-hand side performs the same averaging in two steps: it first chooses the location x according to the intensity measure, and then samples the process according to its Palm distribution $\mathbb{P}^x$. Thus, the Palm kernel formalizes what the surrounding configuration typically looks like around a point of the process.

The Λ_η-almost everywhere uniqueness essentially says that the Palm kernel is only determined at locations where points may appear in expectation. Changing $\mathbb{P}^x$ on a set with Λ_η-measure zero does not affect the relevant identity, so such changes are irrelevant.

In general, these distributions can look very different depending on which x they are taken at. However, if we know that η is stationary on a grouped LC Polish space (S, H), we obtain the following identity via [24, Prop. 3.13]: For any $h \in H$, $z \in S$,

$$\eta^{hz} \overset{d}{=} \theta_h \eta^z. \tag{10}$$

What this means is that the Palm distribution taken at a shifted point is equivalent to the shift of the Palm distribution of the original point, but their ambient structures are identical (up to group action). This conveniently also gives that the Palm kernel is unique on any orbit $Hz := \{h \cdot z \mid h \in H\}$ with $\Lambda_\eta(Hz) > 0$,

since uniqueness at one point of a Λ_η-non-null orbit Hz propagates to all points of that orbit by shifting. This is much more useful than Λ_η-almost everywhere uniqueness, as it allows us to work at any point in a non-null orbit and be sure of the existence of the distribution.

At last, we can bring the conversation back to adjacency processes. Consider a stationary finite adjacency process ξ on a square grouped LC Polish space $(S, H)^2$, and now assume that H acts transitively on S, i.e. that for any $x, y \in S$, one has $y \in Hx$. Then all diagonal points lie in a single orbit, so once a Palm version is fixed at one diagonal point (x, x), the Palm versions at all other diagonal points are determined from it by shifts. In particular, for any $x \in S$ and $h \in H$, Equation (10) gives

$$\xi^{(hx,hx)} =^d \theta_h \xi^{(x,x)}.$$

To pass from this to the constructed graph, note that the diagonal action only relabels the adjacency structure: the map $y \mapsto h \cdot y$ is bijective with inverse $y \mapsto h^{-1} \cdot y$, and each edge atom (y, z) is sent to $(h \cdot y, h \cdot z)$. Thus, $(\theta_h \tau, hx)$ and (τ, x) define isomorphic rooted graphs for every regular rooted adjacency measure (τ, x). Since G is measurable by Lemma 3.5, we may apply it to this distributional identity. Hence, for any $h \in H$,

$$G(\xi^{(x,x)}, x) \stackrel{d}{=} G(\xi^{(hx,hx)}, hx) \tag{11}$$

and therefore the graphs constructed from these Palm distributions are equivalent in law for all diagonal roots. This only leads to one natural question: Is this equivalence of spatial roots enough to conclude that

$$G(\xi^{(x,x)}, x) \stackrel{d}{=} G(\xi, U),$$

where $U \sim \mathbb{U}(\xi, \cdot)$? Unfortunately, this does not yet hold. We first give a counterexample to explain why, and then provide the correct construction. Consider the following finite regular adjacency process on $[-1, 1)^2$. Let A, B be independent uniform random variables on $[-1, 1)$, and define

$$\xi \stackrel{d}{=} \begin{cases} \delta_{A,A} & \text{with probability } 1/2, \\ \delta_{A,A} + \delta_{B,B} + \delta_{A,B} + \delta_{B,A} & \text{with probability } 1/2. \end{cases}$$

This is stationary with respect to the diagonal action of the torus translation group $H = [-1, 1)$ that acts on S transitively via

$$h \cdot x := \begin{cases} x + h - 2 & x + h \geq 1, \\ x + h + 2 & x + h < -1, \\ x + h & -1 \leq x + h < 1. \end{cases}$$

Uniform rooting first samples ξ, and then chooses a vertex uniformly from $V(\xi)$. Hence, letting $U \sim \mathbb{U}(\xi, \cdot)$, we have that

$$G(\xi, U) \stackrel{d}{=} \begin{cases} K_1 & \text{with probability } 1/2, \\ K_2 & \text{with probability } 1/2. \end{cases}$$

By contrast, the Palm distribution at a diagonal point samples the process from a typical diagonal atom. Suppose we take the Palm root at $(0,0)$. Then there are three equally weighted ways for this diagonal atom to arise: the single vertex outcome contributes the atom (A,A), while the two vertex outcome contributes the two possible diagonal atoms (A,A) and (B,B). After fixing the selected diagonal atom at $(0,0)$, the one vertex case gives only the rooted diagonal atom, whereas the two vertex case leaves one additional uniformly distributed vertex C. Thus, the Palm distribution[9] gives,

$$\xi^{(0,0)} \overset{d}{=} \begin{cases} \delta_{0,0} & \text{with probability } 1/3, \\ \delta_{0,0} + \delta_{C,C} + \delta_{0,C} + \delta_{C,0} & \text{with probability } 2/3, \end{cases}$$

and consequently

$$G(\xi^{(0,0)}, 0) \overset{d}{=} \begin{cases} K_1 & \text{with probability } 1/3, \\ K_2 & \text{with probability } 2/3. \end{cases}$$

Evidently, $G(\xi, U) \neq^d G(\xi^{(0,0)}, 0)$. In order to repair this, we need to use a variant of the Palm distribution that counteracts this bias. See the following.

Definition 19 Vertex-Normalized Palm Kernel

Let S^2 be a square LC Polish space, and let ξ be a finite adjacency process on S^2 with intensity measure Λ_ξ. A *Vertex-Normalized Palm kernel* of ξ is a Markov kernel

$$\hat{\mathbb{P}} : S^2 \times \mathcal{B}(\mathcal{E}(S^2)) \to [0,1],$$

written as

$$\hat{\mathbb{P}}^z(A) := \hat{\mathbb{P}}(z, A),$$

such that for every measurable $f : S^2 \times \mathcal{E}(S^2) \to [0,\infty)$,

$$\mathbb{E}\left[\sum_{y \in \xi} \frac{f(y,\xi)}{\xi(\Delta^2(S))}\right] = \int_{S^2} \hat{\mathbb{E}}^z[f(z,\xi)]\mathbb{E}^z\left[\frac{1}{\xi(\Delta^2(S))}\right] \Lambda_\xi(dz),$$

where

$$\hat{\mathbb{E}}^x[f(x,\xi)] := \int_{\mathcal{E}(S^2)} f(z,\tau)\, \hat{\mathbb{P}}^z(d\tau).$$

Here $\Delta^2(S)$ denotes the set $\{(x,x) \mid x \in S\}$. This kernel is equivalent by [26, Eq. (4.8)] to taking the classical Palm distribution at a normalized version of ξ, and therefore again exists and is unique Λ_ξ-almost everywhere, and is subject to the same benefits of stationarity and transitivity as the classical Palm distribution. Analogously to before, whenever $\hat{\xi}^z$ is a point process with law $\hat{\mathbb{P}}^z$, we write

$$\hat{\xi}^z \sim \hat{\mathbb{P}}^z(\cdot)$$

[9] See the formal derivation of this in Appendix B.1.

and refer to $\hat{\xi}^z$ as a *vertex-normalized Palm version* of ξ at z. One can then again view this as a transformation mapping $(\xi, z) \mapsto \hat{\xi}^z$, which we refer to as the *vertex-normalized Palm transformation* of ξ henceforth. Applied in the prior example, one derives[10] instead

$$\hat{\xi}^{(0,0)} \stackrel{d}{=} \begin{cases} \delta_{0,0} & \text{with probability } 1/2, \\ \delta_{0,0} + \delta_{C,C} + \delta_{0,C} + \delta_{C,0} & \text{with probability } 1/2, \end{cases}$$

and consequently

$$G(\hat{\xi}^{(0,0)}, 0) \stackrel{d}{=} \begin{cases} K_1 & \text{with probability } 1/2, \\ K_2 & \text{with probability } 1/2. \end{cases}$$

In this case, $G(\xi, U) =^d G(\hat{\xi}^{(0,0)}, 0)$. In fact, it can be shown that this holds in general, finally giving us a suitable transformation from which to apply Theorem 4.1. This proposition is given below.

▶ **Proposition 5.1.** *Let ξ be a stationary finite adjacency process on a square grouped LC Polish space $(S, H)^2$. Assume that H acts transitively on S. Then, for every $(x, x) \in \Delta^2(S)$,*

$$G(\xi, U) \stackrel{d}{=} G(\hat{\xi}^{(x,x)}, x),$$

where $U \sim \mathbb{U}(\xi, \cdot)$.

Fortunately, to apply the results, one does not (in general) need to formally go through this alternative Palm construction. Instead, we provide a checkable condition under which vertex-normalized Palm variants and classical Palm variants behave the same asymptotically. Since they share a limit in this case, computing the Palm limit is sufficient to also obtain the vertex-normalized Palm limit, which can then be used to construct a local weak limiting graph. This result cannot yet be stated, as we first need to formalize how to work with stationary adjacency processes in sequences.

5.2 α-Admissible Sequences of Spaces

We now discuss how to apply the above results to sequences of finite adjacency processes. There is a significant issue that we need to tackle, which is that it is not possible to have a finite stationary point process on an LC Polish space under our conditions, unless the space is entirely compact. More specifically, if S is a non-compact LC Polish space, there is no grouped LC Polish space (S, H) with H acting transitively on S that admits finite stationary point processes, and thus there is no corresponding square grouped LC Polish space $(S, H)^2$ on which we can have well-behaved finite adjacency processes.[11]

[10] Again see Appendix B.1 for the full derivation.

[11] The proof of this claim is out of scope and irrelevant here beyond motivation.

A perfect example of this is the collection of adjacency processes on $\mathbb{R}^2$ with group action given by the diagonal action of the translation group $(\mathbb{R}, +)$. This space is LC Polish but is not compact, and intuitively, there is no way to have an adjacency process be stationary while also being finite, as any atom that appears at one location must be just as likely to appear after shifting both coordinates by any real amount. Thus, the mass cannot be concentrated in any bounded region. A finite process would have to place only finitely many atoms in this unbounded space, which is incompatible with this translation symmetry. The only possible compatible variant is the degenerate point process that is always empty, which we recall is not included in our use of the term finite.

Thus, the correct finite approximation setting is not stationarity on the limiting space S itself, but stationarity on a sequence of finite or compact subspaces S_n, each of which have a compatible group action. If we take our sequence of finite adjacency processes to be stationary on this sequence of compact subspaces, then we can apply Proposition 5.1 to each element of the sequence, and then consider the convergence of the Palm variants within the ambient space. To ensure that this is possible, we give the following definition of an admissible sequence of spaces, which we use in later results.

Definition 20 α-Admissible Sequence of Spaces

Let S^2 be a square LC Polish space, and fix some $x \in S$. Let $(S_n, H_n)^2_{n\in\mathbb{N}}$ be a sequence of square grouped LC Polish spaces such that for all $n \in \mathbb{N}$, it holds that

1. $x \in S_n$;
2. $S_n \subset S$ elementwise (though with possibly different metric and/or topology);
3. S_n is compact with respect to its topology;
4. The identity map $\mathrm{id}_{S_n} : S_n \to S$ is measurable with respect to the Borel σ-algebras induced by the topology on S_n and the topology on S;
5. H_n acts transitively on S_n for every $n \in \mathbb{N}$.

Then, we call $(S_n, H_n)^2_{n\in\mathbb{N}}$ an *α-admissible sequence of spaces* in S^2 sharing $(x, x) \in S^2$, and denote it as $(S_n, H_n)^2_{n\in\mathbb{N}} \subset_{\alpha,x} S^2$.

The use of α here has no semantic meaning, and is just used to disambiguate this class of sequences of spaces from a more restricted class with stronger conditions, the class of β-admissible sequences of spaces, which is introduced later to handle specific model types.

Now a quick comment on the stated conditions. Conditions (1)-(3) encode the basic finite approximation setting: the approximating spaces must sit inside the ambient space as sets, remain compact in their intrinsic topology, and share the reference point x at which the Palm variants are taken. Condition (5) gives transitivity on each S_n, which is exactly what allows the Palm identities from

the previous subsection to be transferred across all diagonal points. Condition (4) is more technical. Its role is to ensure that an adjacency process defined on S_n can also be regarded canonically as an S-valued process by pushing it forward along the identity inclusion $S_n \to S$. The subspace view is what is used for stationarity and Palm theory on $(S_n, H_n)^2$, while the ambient view is what is used for weak convergence in $\mathcal{E}^\circ(S^2)$. When we have such an adjacency process on a subspace S_n, we adopt the convention of denoting its pushforward to S under the identity inclusion map via the same symbol used to denote the subspace process, to avoid introducing unnecessary notation.

The classical example of such a family, used in [33] and analogous to in [18], is the α-admissible sequence of spaces $(\mathbb{T}_n^d, \mathbb{T}_n^d)_{n\in\mathbb{N}}^2 \subset_{\alpha,0} \mathbb{R}^{2d}$ where $\mathbb{T}_n^d$ is the flat n-volume torus in $\mathbb{R}^d$

$$\mathbb{T}_n^d := \mathbb{R}^d/n^{1/d}\mathbb{Z}^d = \left[-\frac{n^{1/d}}{2}, \frac{n^{1/d}}{2}\right)^d$$

acting on itself via the action

$$h \cdot x := x + h \pmod{n^{1/d}\mathbb{Z}^d}.$$

Indeed, these spaces are all compact with respect to their local topology, namely the quotient topology, they all share $(0, \ldots, 0) \in \mathbb{R}^{2d}$, and the actions are transitive. There are several more exotic examples that are also α-admissible. Some examples are given below,

1. Complex flat tori $S_n := \mathbb{C}^d/n^{1/d}(\mathbb{Z}^d + i\mathbb{Z}^d) \subset \mathbb{C}^d$, with $H_n := S_n$ acting transitively on themselves by addition, though the metric is analogous to that of $\mathbb{R}^{2d}$, and thus this does not provide much of interest for isometric models.
2. p-adic balls $S_n := p^{-n}\mathbb{Z}_p \subset \mathbb{Q}_p$, with $H_n := (p^{-n}\mathbb{Z}_p, +)$ acting transitively by translation.
3. Discrete tori $S_n := \mathbb{Z}^d/n\mathbb{Z}^d \subset \mathbb{Z}^d$, with $H_n := S_n$ acting transitively on itself by addition.
4. Shifted scaled unitary groups $S_n := I + n(\mathrm{U}(d) - I) \subset M_d(\mathbb{C})$, with the translated left-right action of $H_n := \mathrm{U}(d) \times \mathrm{U}(d)$.

There are, of course, many other conceivable examples. Finally, we have the necessary context to give the local weak convergence theorem.

5.3 Local Weak Convergence Theorem

As before, we will give the statement of the theorem, then a proof of this theorem dependent on intermediate results, and then the proof of these dependencies in the later sections.

▶ **Theorem 5.2** Local Weak Convergence Theorem. *Let S^2 be a square LC Polish space and fix $x \in S$. Let $(\xi_n)_{n\in\mathbb{N}}$ be a sequence of finite stationary adjacency processes on an α-admissible sequence of spaces $(S_n, H_n)^2 \subset_{\alpha,x} S^2$. Then, if it holds that*

1. $\xi_n^{(x,x)} \Rightarrow \xi_\infty$ *in* $\mathcal{E}^\circ(S^2)$ *for some regular adjacency process* ξ_∞,
2. $(\xi_n^{(x,x)}, x)$ *is a radially tight sequence,*
3. $$\mathbb{E}\left[\left|\frac{|V(\xi_n)|}{\mathbb{E}[|V(\xi_n)|]} - 1\right|\right] \to 0,$$

it also holds that

$$G(\xi_n, U_n) \Rightarrow G(\xi_\infty, x),$$

where $U_n \sim \mathbb{U}(\xi_n, \cdot)$.

Essentially, this says that if we can show that the Palm variants converge, and that the vertex counts concentrate in a suitable way, we obtain local weak convergence. There are only two remaining lemmas to state. The first gives the asymptotic equivalence of the classical and vertex-normalized Palm variants under this vertex concentration, while the second guarantees that radial tightness passes between these two Palm variants.

First, we have the asymptotic equivalence lemma. This uses the notion of convergence in total variation of the difference between respective probability laws, which formally is defined to be

$$\left\|\hat{\mathbb{P}}_n^{(x,x)} - \mathbb{P}_n^{(x,x)}\right\|_{\mathrm{TV}} \to 0 \tag{12}$$

where

$$\left\|\hat{\mathbb{P}}_n^{(x,x)} - \mathbb{P}_n^{(x,x)}\right\|_{\mathrm{TV}} := \sup_{A \in \mathcal{B}(\mathcal{E}(S^2))} \left|\hat{\mathbb{P}}_n^{(x,x)}(A) - \mathbb{P}_n^{(x,x)}(A)\right|.$$

Essentially, this enforces the condition that the probability of any event occurring gets arbitrarily close as n gets large. This convergence guarantees that the two sequences $\xi_n^{(x,x)}$ and $\hat{\xi}_n^{(x,x)}$ share a limit, i.e. that for any ξ_∞,

$$\xi_n^{(x,x)} \Rightarrow \xi_\infty \quad \Longleftrightarrow \quad \hat{\xi}_n^{(x,x)} \Rightarrow \xi_\infty. \tag{13}$$

This is given by [38, Prop 18.2] along with the triangle inequality on the TV norm. Thus, we give the lemma along with a usable corollary.

▶ **Lemma 5.3.** *Let* S^2 *be a square LC Polish space and fix* $x \in S$. *Let* $(\xi_n)_{n\in\mathbb{N}}$ *be a sequence of finite stationary adjacency processes on an* α*-admissible sequence of spaces* $(S_n, H_n)^2 \subset_{\alpha,x} S^2$. *Then, if it holds that*

$$\mathbb{E}\left[\left|\frac{|V(\xi_n)|}{\mathbb{E}[|V(\xi_n)|]} - 1\right|\right] \to 0,$$

it also holds that

$$\left\|\hat{\mathbb{P}}_n^{(x,x)} - \mathbb{P}_n^{(x,x)}\right\|_{\mathrm{TV}} \longrightarrow 0.$$

The vertex concentration condition here essentially says that $|V(\xi_n)|$ becomes close to its expectation relative to the scale of that expectation. This matters

because the classical Palm law is biased toward configurations with more diagonal atoms, while the vertex-normalized Palm law corrects for this bias. When the vertex count concentrates, this correction becomes asymptotically negligible, yielding the total variation convergence in Equation (12). The corollary below follows immediately.

▶ **Corollary 5.4.** *Let S^2 be a square LC Polish space and fix $x \in S$. Let $(\xi_n)_{n\in\mathbb{N}}$ be a sequence of finite stationary adjacency processes on an α-admissible sequence of spaces $(S_n, H_n)^2 \subset_{\alpha,x} S^2$. Then, if it holds that*

$$\mathbb{E}\left[\left|\frac{|V(\xi_n)|}{\mathbb{E}[|V(\xi_n)|]} - 1\right|\right] \to 0,$$

it also holds that for any $\xi_\infty \in \mathcal{E}(S^2)$,

$$\xi_n^{(x,x)} \Rightarrow \xi_\infty \quad \Longleftrightarrow \quad \hat{\xi}_n^{(x,x)} \Rightarrow \xi_\infty.$$

Proof. By Lemma 5.3, we have

$$\left\|\hat{\mathbb{P}}_n^{(x,x)} - \mathbb{P}_n^{(x,x)}\right\|_{\mathrm{TV}} \to 0.$$

The claimed equivalence then follows immediately from Equation (13). ◀

It also is necessary to show that radial tightness translates between these two variants. Specifically,

▶ **Lemma 5.5.** *Let S^2 be a square LC Polish space and fix $x \in S$. Let $(\xi_n)_{n\in\mathbb{N}}$ be a sequence of finite stationary adjacency processes on an α-admissible sequence of spaces $(S_n, H_n)^2 \subset_{\alpha,x} S^2$. Suppose further that*

$$\mathbb{E}\left[\left|\frac{|V(\xi_n)|}{\mathbb{E}[|V(\xi_n)|]} - 1\right|\right] \to 0.$$

Then, it holds that $(\xi_n^{(x,x)}, x)_{n\in\mathbb{N}}$ is radially tight if and only if $(\hat{\xi}_n^{(x,x)}, x)_{n\in\mathbb{N}}$ is radially tight.

With these lemmas, and with the priorly stated Proposition 5.1, we can give the proof of Theorem 5.2 conditional on the proofs of these results.

Proof of the Local Weak Convergence Theorem (Theorem 5.2)

Proof. We want to show that for every $r \in \mathbb{N}$ and $\tilde{G} \in \mathcal{G}_*$, it holds that

$$\mathbb{P}(B_r(\xi_n, U_n) \cong \tilde{G}) \to \mathbb{P}(B_r(\xi_\infty, x) \cong \tilde{G})$$

Fix such an $r \in \mathbb{N}$ and $\tilde{G} \in \mathcal{G}_*$. By the properties of α-admissible sequences of spaces, Definition 20, we know that both

$$(\xi_n^{(x,x)}, x)$$

and

$$(\hat{\xi}_n^{(x,x)}, x)$$

exist uniquely for all $n \in \mathbb{N}$, and by assumption we furthermore have that

$$(\xi_n^{(x,x)}, x) \Rightarrow (\xi_\infty, x).$$

By Corollary 5.4, it also holds that

$$(\hat{\xi}_n^{(x,x)}, x) \Rightarrow (\xi_\infty, x). \tag{14}$$

Since we have radial tightness of $(\xi_n^{(x,x)}, x)_{n\in\mathbb{N}}$ by assumption, we also have radial tightness of $(\hat{\xi}_n^{(x,x)}, x)_{n\in\mathbb{N}}$ via Lemma 5.5. By Proposition 5.1, we have that

$$\mathbb{P}(B_r(\xi_n, U_n) \cong \tilde{G}) = \mathbb{P}(B_r(\hat{\xi}_n^{(x,x)}, x) \cong \tilde{G}),$$

and it therefore suffices to show that

$$\mathbb{P}(B_r(\hat{\xi}_n^{(x,x)}, x) \cong \tilde{G}) \to \mathbb{P}(B_r(\xi_\infty, x) \cong \tilde{G}). \tag{15}$$

By Theorem 4.1 with Equation (14) and the radial tightness of the vertex-normalized Palm sequence, we have

$$G(\hat{\xi}_n^{(x,x)}, x) \Rightarrow G(\xi_\infty, x),$$

which gives Equation (15), concluding the proof. ◀

5.4 Propositions

There is only one proposition that we need to prove, being Proposition 5.1. The idea of the proof is to rewrite uniform rooting as averaging over diagonal atoms, then apply the vertex-normalized Palm identity to transfer that average to Palm space. Transitivity makes the rooted-ball law independent of which diagonal point is used, so one factor becomes constant, and a final normalization identity shows the remaining factor is one. This gives equality in distribution between the uniformly rooted graph and the graph rooted by the vertex-normalized Palm version.

Proof of Proposition 5.1

Proof. Fix $(x, x) \in \Delta^2(S)$. It suffices to show that for every $r \in \mathbb{N}$ and every finite rooted graph $\tilde{G} \in \mathcal{G}_*$,

$$\mathbb{P}\Big(B_r(\xi, U) \cong \tilde{G}\Big) = \mathbb{P}\Big(B_r(\hat{\xi}^{(x,x)}, x) \cong \tilde{G}\Big). \tag{16}$$

Fix such $r \in \mathbb{N}$ and $\tilde{G} \in \mathcal{G}_*$. Define $f : S^2 \times \mathcal{E}(S^2) \to \{0, 1\}$ by

$$f(z, \tau) := \mathbf{1}_{\{z\in\Delta^2(S)\}}\mathbf{1}_{\{B_r(\tau,\pi_1(z))\cong\tilde{G}\}}, \qquad z \in S^2,\ \tau \in \mathcal{E}(S^2),$$

where $\pi_1(x,y) = x$. This is measurable by measurability of $z \mapsto \pi_1(z)$ and $B_r(\cdot)$. Since $U \sim \mathbb{U}(\xi, \cdot)$, we have that

$$\mathbb{P}\Big(B_r(\xi, U) \cong \tilde{G}\Big) = \mathbb{E}\left[\frac{1}{|V(\xi)|} \sum_{y \in V(\xi)} \mathbf{1}_{\{B_r(\xi,y) \cong \tilde{G}\}}\right].$$

Because ξ is a finite adjacency process, $\xi(\Delta^2(S)) = |V(\xi)| < \infty$ almost surely, hence

$$\mathbb{P}\Big(B_r(\xi, U) \cong \tilde{G}\Big) = \mathbb{E}\left[\sum_{z \in \xi} \frac{f(z,\xi)}{\xi(\Delta^2(S))}\right].$$

Applying the vertex-normalized Palm identity from Definition 19 yields

$$\mathbb{P}\Big(B_r(\xi, U) \cong \tilde{G}\Big) = \int_{S^2} \hat{\mathbb{E}}^z[f(z,\xi)]\, \mathbb{E}^z\left[\frac{1}{\xi(\Delta^2(S))}\right] \Lambda_\xi(dz).$$

Since $f(z,\tau) = 0$ for $z \notin \Delta^2(S)$, this reduces to

$$\mathbb{P}\Big(B_r(\xi, U) \cong \tilde{G}\Big) = \int_{\Delta^2(S)} \hat{\mathbb{E}}^z[f(z,\xi)]\, \mathbb{E}^z\left[\frac{1}{\xi(\Delta^2(S))}\right] \Lambda_\xi(dz).$$

We now identify the first factor. Let $z \in \Delta^2(S)$. By transitivity, there exists $h \in H$ with $z = h(x,x)$. Using the covariance identity in random-variable form for the vertex-normalized Palm versions, Equation (10), and Definition 19,

$$\hat{\xi}^z \overset{d}{=} \theta_h \hat{\xi}^{(x,x)}.$$

The diagonal action preserves rooted r-balls since it acts as a relabeling of the vertices, and thus analogously to Equation (11),

$$B_r(\theta_h \tau, hx) \cong B_r(\tau, x).$$

Therefore, this factor is some constant,

$$\hat{\mathbb{E}}^z[f(z,\xi)] = \mathbb{P}\Big(B_r(\hat{\xi}^{(x,x)}, x) \cong \tilde{G}\Big) =: c_{r,\tilde{G}},$$

which is independent of $z \in \Delta^2(S)$. Hence,

$$\mathbb{P}\Big(B_r(\xi, U) \cong \tilde{G}\Big) = c_{r,\tilde{G}} \int_{\Delta^2(S)} \mathbb{E}^z\left[\frac{1}{\xi(\Delta^2(S))}\right] \Lambda_\xi(dz). \tag{17}$$

It remains to compute the integral. Apply the same vertex-normalized Palm identity from Definition 19 with

$$g(z,\tau) := \mathbf{1}_{\{z \in \Delta^2(S)\}}.$$

Then,

$$\begin{aligned}\int_{\Delta^2(S)} \mathbb{E}^z\left[\frac{1}{\xi(\Delta^2(S))}\right]\Lambda_\xi(dz) &= \int_{S^2} \mathbf{1}_{\{z\in\Delta^2(S)\}}\mathbb{E}^z\left[\frac{1}{\xi(\Delta^2(S))}\right]\Lambda_\xi(dz)\\ &= \int_{S^2} \hat{\mathbb{E}}^z\left[\mathbf{1}_{\{z\in\Delta^2(S)\}}\right]\mathbb{E}^z\left[\frac{1}{\xi(\Delta^2(S))}\right]\Lambda_\xi(dz)\\ &= \mathbb{E}\left[\sum_{u\in\xi}\frac{\mathbf{1}_{\{u\in\Delta^2(S)\}}}{\xi(\Delta^2(S))}\right].\end{aligned}$$

The numerator equals $\xi(\Delta^2(S))$, and this denominator is strictly positive almost surely. Therefore,

$$\int_{\Delta^2(S)} \mathbb{E}^z\left[\frac{1}{\xi(\Delta^2(S))}\right]\Lambda_\xi(dz) = 1. \tag{18}$$

Substituting Equation (18) into Equation (17) gives

$$\mathbb{P}\Big(B_r(\xi,U)\cong\tilde{G}\Big) = \mathbb{P}\Big(B_r(\hat{\xi}^{(x,x)},x)\cong\tilde{G}\Big).$$

This is Equation (16). Since it holds for all $r\in\mathbb{N}$ and all finite rooted test graphs, we conclude

$$G(\xi,U)\overset{d}{=}G(\hat{\xi}^{(x,x)},x).$$

◀

5.5 Lemmas

We begin with Lemma 5.3, as its result is necessary for the subsequent Lemma 5.5. In order to proceed with the proof, we first need one additional lemma, giving an identity that relates the classical and vertex-normalized Palm distributions.

▶ **Lemma 5.6.** *Let S^2 be a square LC Polish space and fix $x\in S$. Take some finite adjacency process ξ taking values in $\mathcal{E}(S^2)$. Define*

$$w(\tau) := \frac{1}{\tau(\Delta^2(S))}, \qquad \tau\in\mathcal{E}(S^2).$$

Let $\mathbb{P}^{(x,x)}$ and $\hat{\mathbb{P}}^{(x,x)}$ be some corresponding Palm and vertex-normalized Palm variants that exist uniquely for some $(x,x)\in S$. Then, it holds that

$$\hat{\mathbb{P}}^{(x,x)}(A) = \frac{\mathbb{E}^{(x,x)}[\mathbf{1}_A(\xi)\,w(\xi)]}{\mathbb{E}^{(x,x)}[w(\xi)]}$$

for all Borel $A\subseteq\mathcal{E}(S^2)$.

Proof. Define a kernel

$$\widetilde{\mathbb{P}}^{\,z}(A) := \frac{\mathbb{E}^z[\mathbf{1}_A(\xi)\, w(\xi)]}{\mathbb{E}^z[w(\xi)]}, \qquad z \in S^2.$$

Let $\widetilde{\mathbb{E}}^z$ be the expectation under $\widetilde{\mathbb{P}}^{\,z}$. For any measurable $f : S^2 \times \mathcal{E}(S^2) \to [0,\infty)$,

$$\begin{aligned}\int_{S^2} \widetilde{\mathbb{E}}^z[f(z,\xi)]\, \mathbb{E}^z[w(\xi)]\, \Lambda_\xi(dz) &= \int_{S^2} \mathbb{E}^z[f(z,\xi)w(\xi)]\, \Lambda_\xi(dz) \\ &= \mathbb{E}\left[\sum_{u\in\xi} f(u,\xi)w(\xi)\right] \\ &= \mathbb{E}\left[\sum_{u\in\xi} \frac{f(u,\xi)}{\xi(\Delta^2(S))}\right].\end{aligned}$$

The second line is the ordinary Palm identity (Definition 18), and the last line is $w = 1/\xi(\Delta^2(S))$. Hence, $\widetilde{\mathbb{P}}$ satisfies exactly the defining identity of a vertex-normalized Palm kernel (Definition 19). By uniqueness of that kernel for the chosen versions at (x,x), $\widetilde{\mathbb{P}}^{(x,x)} = \hat{\mathbb{P}}^{(x,x)}$, which is the claim. ◀

Of course this extends to all (x,x) when we have uniqueness of the diagonal via stationarity and diagonal transitivity. Now the proof of Lemma 5.3. The proof rewrites the vertex-normalized Palm law as a weighted version of the ordinary Palm law using the above lemma, and then bounds their total-variation distance by the fluctuation of that weight under the ordinary Palm law. Using the Palm identity alongside stationarity and transitivity, these Palm expectations are converted to original-process expectations, and the bound reduces exactly to the assumed vertex-count concentration term. Since that term tends to zero, the total-variation distance also tends to zero.

Proof of Lemma 5.3

Proof. For each n, define

$$m_n := \mathbb{E}[|V(\xi_n)|] \qquad \text{and} \qquad w(\tau) := \frac{1}{|V(\tau)|} = \frac{1}{\tau(\Delta^2(S))}.$$

Because ξ_n is finite (hence non-empty via our convention), we have

$$1 \le |V(\xi_n)| < \infty \quad \text{almost surely,}$$

so $0 < w_n(\xi_n) \le 1$ almost surely and $0 < m_n < \infty$. By Lemma 5.6,

$$\hat{\mathbb{P}}_n^{(x,x)}(A) = \frac{\mathbb{E}_n^{(x,x)}[\mathbf{1}_A(\xi_n)\, w(\xi_n)]}{\mathbb{E}_n^{(x,x)}[w(\xi_n)]}$$

for all Borel $A \subseteq \mathcal{E}(S^2)$. Therefore,

$$\begin{aligned}\hat{\mathbb{P}}_n^{(x,x)}(A) - \mathbb{P}_n^{(x,x)}(A) &= \frac{\mathbb{E}_n^{(x,x)}[\mathbf{1}_A(\xi_n)\, w(\xi_n)]}{\mathbb{E}_n^{(x,x)}[w(\xi_n)]} - \mathbb{E}_n^{(x,x)}[\mathbf{1}_A(\xi_n)] \\ &= \frac{\mathbb{E}_n^{(x,x)}[\mathbf{1}_A(\xi_n)\, w(\xi_n)] - \mathbb{E}_n^{(x,x)}[\mathbf{1}_A(\xi_n)]\, \mathbb{E}_n^{(x,x)}[w(\xi_n)]}{\mathbb{E}_n^{(x,x)}[w(\xi_n)]} \\ &= \frac{\mathbb{E}_n^{(x,x)}\Big[\mathbf{1}_A(\xi_n)\Big(w(\xi_n) - \mathbb{E}_n^{(x,x)}[w(\xi_n)]\Big)\Big]}{\mathbb{E}_n^{(x,x)}[w(\xi_n)]}.\end{aligned}$$

Taking absolute values and then $\sup_A$ gives, letting $w_n := w(\xi_n)$,

$$\big\|\hat{\mathbb{P}}_n^{(x,x)} - \mathbb{P}_n^{(x,x)}\big\|_{\mathrm{TV}} \le \frac{\mathbb{E}_n^{(x,x)}\Big[\Big|w_n - \mathbb{E}_n^{(x,x)}[w_n]\Big|\Big]}{\mathbb{E}_n^{(x,x)}[w_n]}. \tag{19}$$

It remains to compute the two Palm expectations in the bound, which we do now. Let $G : \mathcal{E}(S^2) \to [0,\infty)$ be measurable and invariant under diagonal shifts, i.e. $G(\theta_h \tau) = G(\tau)$ for all $h \in H_n$. Applying the ordinary Palm identity with $f(z,\tau) := \mathbf{1}_{\{z \in \Delta^2(S)\}} G(\tau)$, we obtain

$$\mathbb{E}[|V(\xi_n)|\, G(\xi_n)] = \int_{\Delta^2(S)} \mathbb{E}_n^z[G(\xi_n)]\, \Lambda_{\xi_n}(dz).$$

Fix $z \in \Delta^2(S)$. By transitivity on S_n, there is $h_z \in H_n$ with $z = h_z(x,x)$. Using the covariance relation in random-variable form from Equation (10),

$$\xi_n^z \overset{d}{=} \theta_{h_z} \xi_n^{(x,x)},$$

and H_n-invariance of G, we get

$$\mathbb{E}_n^z[G(\xi_n)] = \mathbb{E}[G(\xi_n^z)] = \mathbb{E}[G(\theta_{h_z}\xi_n^{(x,x)})] = \mathbb{E}[G(\xi_n^{(x,x)})] = \mathbb{E}_n^{(x,x)}[G(\xi_n)].$$

Thus,

$$\mathbb{E}[|V(\xi_n)|\, G(\xi_n)] = \Lambda_{\xi_n}(\Delta^2(S))\, \mathbb{E}_n^{(x,x)}[G(\xi_n)].$$

Taking $G \equiv 1$ yields

$$\Lambda_{\xi_n}(\Delta^2(S)) = \mathbb{E}[|V(\xi_n)|] = m_n,$$

hence for every non-negative diagonal-shift-invariant G,

$$\mathbb{E}_n^{(x,x)}[G(\xi_n)] = \frac{\mathbb{E}[|V(\xi_n)|\, G(\xi_n)]}{m_n}.$$

Apply this first with $G = w$, which is diagonal-shift-invariant because diagonal actions relabel vertices bijectively and preserve $|V(\cdot)|$. Then

$$\mathbb{E}_n^{(x,x)}[w_n] = \frac{\mathbb{E}\big[|V(\xi_n)|\, |V(\xi_n)|^{-1}\big]}{m_n} = \frac{1}{m_n}.$$

Next apply it with

$$G(\tau) := \left| w(\tau) - \frac{1}{m_n} \right|,$$

which is also diagonal-shift-invariant. This gives

$$\mathbb{E}_n^{(x,x)}\left[\left| w_n - \frac{1}{m_n} \right|\right] = \frac{1}{m_n}\mathbb{E}\left[|V(\xi_n)|\left|\frac{1}{|V(\xi_n)|} - \frac{1}{m_n}\right|\right] = \frac{1}{m_n}\mathbb{E}\left[\left|1 - \frac{|V(\xi_n)|}{m_n}\right|\right].$$

Substitute these into the total-variation bound, Equation (19),

$$\left\|\hat{\mathbb{P}}_n^{(x,x)} - \mathbb{P}_n^{(x,x)}\right\|_{\mathrm{TV}} \le \frac{\frac{1}{m_n}\mathbb{E}\left[\left|1 - \frac{|V(\xi_n)|}{m_n}\right|\right]}{\frac{1}{m_n}} = \mathbb{E}\left[\left|\frac{|V(\xi_n)|}{\mathbb{E}[|V(\xi_n)|]} - 1\right|\right].$$

By hypothesis, the right-hand side tends to 0, so

$$\left\|\hat{\mathbb{P}}_n^{(x,x)} - \mathbb{P}_n^{(x,x)}\right\|_{\mathrm{TV}} \longrightarrow 0.$$

◀

Finally, we give the proof of Lemma 5.5. The strategy is to bound the gap between the two failure probabilities by their total variation distance, then use a compact containment bound for one Palm variant and the vanishing gap to transfer the same bound to the other.

Proof of Lemma 5.5

Proof. Fix $r \in \mathbb{N}$, a compact $L \subset S$, and define

$$A_{r,L} := \{\tau \in \mathcal{E}(S^2) : V_r(\tau, x) \not\subset L\}.$$

Since V_r is measurable by Lemma 3.6, $A_{r,L}$ is measurable.

For each n, define

$$\delta_n := \left\|\hat{\mathbb{P}}_n^{(x,x)} - \mathbb{P}_n^{(x,x)}\right\|_{\mathrm{TV}},$$

using the total-variation definition in Equation (12). By Lemma 5.3, $\delta_n \to 0$. By the definition of total variation, for every measurable set A,

$$\hat{\mathbb{P}}_n^{(x,x)}(A) \le \mathbb{P}_n^{(x,x)}(A) + \delta_n, \qquad \mathbb{P}_n^{(x,x)}(A) \le \hat{\mathbb{P}}_n^{(x,x)}(A) + \delta_n.$$

Applying this with $A = A_{r,L}$ gives, for each n,

$$\mathbb{P}\Big(V_r(\hat{\xi}_n^{(x,x)}, x) \not\subset L\Big) \le \mathbb{P}\Big(V_r(\xi_n^{(x,x)}, x) \not\subset L\Big) + \delta_n, \tag{20}$$

and

$$\mathbb{P}\Big(V_r(\xi_n^{(x,x)}, x) \not\subset L\Big) \le \mathbb{P}\Big(V_r(\hat{\xi}_n^{(x,x)}, x) \not\subset L\Big) + \delta_n. \tag{21}$$

($\Rightarrow$) Assume $(\xi_n^{(x,x)}, x)_{n\in\mathbb{N}}$ is radially tight. Fix $\varepsilon > 0$. By radial tightness, there exist a compact $L \subset S$ and $N_0 \in \mathbb{N}$ such that for all $n \geq N_0$,

$$\mathbb{P}\Big(V_r(\xi_n^{(x,x)}, x) \not\subset L\Big) < \frac{\varepsilon}{2}.$$

Since $\delta_n \to 0$, there exists $N_1 \in \mathbb{N}$ such that for all $n \geq N_1$, $\delta_n < \varepsilon/2$. Thus for all $n \geq N := \max\{N_0, N_1\}$, Equation (20) gives

$$\mathbb{P}\Big(V_r(\hat{\xi}_n^{(x,x)}, x) \not\subset L\Big) < \frac{\varepsilon}{2} + \frac{\varepsilon}{2} = \varepsilon.$$

Hence, $(\hat{\xi}_n^{(x,x)}, x)_{n\in\mathbb{N}}$ is radially tight.

($\Leftarrow$) The argument is identical by symmetry on Equation (21).

Both implications hold, so the two sequences share radial tightness. ◀

6 Random Graph Models

With the proof of Theorem 5.2, it is now possible for one to show local weak convergence of a graph sequence by verifying

1. vertex concentration;
2. convergence of the Palm transformed adjacency processes;
3. fiber-finiteness of the limiting adjacency process;
4. radial tightness of the sequence of Palm transformed adjacency processes.

While vertex concentration is typically straightforward to verify, the other three conditions are highly nontrivial. General checkable conditions for them at the level of abstraction of Section 5 are likely unattainable beyond restatements of the definitions. However, for certain classes of random graph models, it is possible to give some (relatively) easy to check sufficient conditions.

This section is devoted to exactly this, with the goal of specifying a class of random graph models, and proving such sufficient conditions on these models. This class of models is what we call the class of *Spatial Inhomogeneous Random Graphs*, first defined in [18], though we adapt this definition slightly to make it compatible with our framework. In essence, these models take some random collection of vertices, assign marks to these vertices randomly, and then connect pairs of vertices independently with probability dependent only on the distance between the vertices and their marks.

This notion of "connecting vertices randomly" is an extremely prominent approach in the literature. In fact, it is largely universal. Thus far, we have treated adjacency processes as the elementary objects of our theory, where the adjacency structure is given directly according to some distribution and the vertices are derived from that adjacency structure. However, in the standard literature, one often instead first considers some distribution of vertices[12] in some space, and then connects pairs randomly to arrive at a random graph. This viewpoint is not adversarial to us, but is instead very useful, as we will show that every adjacency process can be constructed in such a way.

In this section, we will give the general theory behind this connective approach, specialize it to spatial inhomogeneous random graphs, and then give a local weak convergence theorem for these models that leverages checkable criteria for the conditions of Theorem 5.2 to invoke the result. As in Section 5, we will also define an admissible sequence of spaces on which this theorem can be applied.

6.1 Connection Kernels and Representability

In the aforementioned types of connection-based models, the most common approach is typically enumeration based, i.e. instead of considering a vertex process Γ without a canonical enumeration, they consider a family of N random

[12] Sometimes these are called vertex labels, but we do not distinguish between labels and spatial positions in this way.

variables $(X_i)_{i\leq N}$ representing vertices and connect vertices based on indices. We will give two such examples, the *Erdős–Rényi random graph* [17, Sec. 1.3.1] and the *random geometric graph (RGG)* [18, Sec. 1].

In the Erdős–Rényi random graph model, one starts with an enumerated family of N vertex variables

$$(X_i)_{i\leq N}.$$

In the classical model these variables carry no geometric information, so one may consider them based only on their indices, and thus the vertex set is then identified with the index set $[N] := \{1, \ldots, N\}$. Given a parameter $p \in [0, 1]$, each pair of distinct indices $i, j \in [N]$ is connected independently with probability p. Thus, the corresponding random graph, via adjacency processes, is

$$\xi_{N,p} := \sum_{i=1}^{N} \delta_{(X_i,X_i)} + \sum_{i=1}^{N} \sum_{j=i+1}^{N} E_{ij}(\delta_{(X_i,X_j)} + \delta_{(X_j,X_i)}) \tag{22}$$

where

$$(E_{ij})_{1\leq i<j\leq N}$$

is an independent family of Bernoulli random variables with common parameter $p \in [0, 1]$. In the RGG model, the family of N vertex variables

$$(X_i)_{i\leq N}$$

is taken in some LC Polish space and does carry geometric information. Most commonly, we assume that all of the variables in this family are i.i.d. with respect to some probability measure μ on S, and the corresponding undirected random graph (again as an adjacency process) is given by

$$\xi_{N,r} := \sum_{i=1}^{N} \delta_{(X_i,X_i)} + \sum_{i=1}^{N} \sum_{j=i+1}^{N} \mathbf{1}_{\{d_S(X_i,X_j)\leq r\}}(\delta_{(X_i,X_j)} + \delta_{(X_j,X_i)}) \tag{23}$$

where r is some real number called the connection radius. Unlike the Erdős–Rényi model, the random variables $(X_i)_{i\leq N}$ are not merely labels or trivial vertex marks: they are spatial locations, and the adjacency relation is derived deterministically from these locations by the metric threshold rule.

There are two things to note about the way we stated the examples. Firstly, although the general theory developed earlier allows adjacency processes that need not be symmetric, in the models considered in this section we now choose to work with symmetric constructions. Concretely, each undirected edge is associated with a single unordered pair of vertices, and when that edge occurs we place both orderings jointly in the adjacency process. In the two examples above this is enforced by summing only over pairs with $j > i$, and then placing both orderings.

This convention is useful because it makes the adjacency process reflect true edge occurrence in the underlying graph without any further adjustment. If an

undirected edge can instead be represented by two ordered adjacency points, then the constructed graph contains that edge whenever either one, or both, of those representations occur, which increases the density of the constructed graph when these points are not jointly placed. For example, in the Erdős–Rényi graph, if the two ordered edges (i, j) and (j, i) were sampled independently with common probability p, then the observed density of undirected edges would be $1 - (1 - p)^2 = 2p - p^2$, since the graph contains $\{i, j\}$ as soon as at least one of the two ordered representations appears.

Avoiding this issue via symmetry also keeps later calculations of graph observables, such as degree or cycle distributions, transparent, since these are computed in the undirected graph we actually study, and so it is preferable that the model encode each potential edge jointly rather than through two separate orientations. In particular, this is the natural convention if one wants to preserve the distributions of graph observables in the sense relevant for models such as those studied in [18].

The other noteworthy design choice to review is our choice of terminology. We refer to these objects as random graphs, even though, strictly speaking, they are adjacency processes. This is justified for two reasons. First, each adjacency process canonically induces a unique random graph, so no structural information relevant to the graph model is lost. Second, this terminology aligns with the existing random graph literature and makes clear which class of models is under discussion.

With the intuition behind connection-based models established, we can give the definition of the connection kernel. We use a kernel to connect vertices rather than using edge indicator variables on enumerated vertices as in the two aforementioned examples. This is not only because it makes the theory cleaner and more measure-theoretical, but also because it allows for more reliability under Palm transformations. When the vertices are given by an enumeration of random variables, there is no canonical way to re-enumerate them after taking the Palm version of the vertex process, and thus the Palm vertex distribution cannot be passed back into the random graph model reliably, which creates major challenges for local weak convergence.

So, instead of using enumerations, we simply define the connection kernel as a kernel that, for every fixed collection of vertices, assigns a distribution of adjacency measures (an adjacency process) measurably in such a way that the vertex identities remain unchanged. See the definition below.

Definition 21 Connection Kernel

Let S^2 be a square LC Polish space. Then, a Markov kernel $C : \mathcal{N}(S) \times \mathcal{B}(\mathcal{A}(S^2)) \to [0, 1]$ is called a *connection kernel* iff for every $\gamma \in \mathcal{N}(S)$, it

holds that $C(\gamma, A) = 0$, where

$$A = \{\tau \in \mathcal{A}(S^2) \mid V(\tau) \neq \gamma\}.$$

This is a Borel set because the singleton $\{\gamma\}$ is closed and the map V is measurable as per Lemma 3.1.

Such a construction always gives adjacency measures, not general counting measures, as guaranteed by the domain. The additional condition ensures that the connection kernel only places edges between existing vertices, and does not change the vertex collection. It also gives the useful identity that whenever an adjacency process ξ is given by a connection kernel C applied to a vertex distribution Γ, i.e. $\xi \sim C(\Gamma, \cdot)$, it holds that

$$V(\xi) = \Gamma$$

almost surely. This can also be seen as an equivalent defining property for connection kernels. It turns out that every adjacency process can be represented in this way.

▶ **Lemma 6.1** Representability via Connection Kernels. *Let S^2 be a square LC Polish space, and let $\xi \in \mathcal{A}(S^2)$ be an adjacency process. Then, there exists some connection kernel C for which $\xi \sim C(V(\xi), \cdot)$.*

Proof. Set $\Gamma := V(\xi)$. By Lemma 3.1, Γ is an $\mathcal{N}(S)$-valued random element. Since $\mathcal{A}(S^2)$ and $\mathcal{N}(S)$ are Polish, there is a regular conditional law K of ξ given Γ, so $\xi \sim K(\Gamma, \cdot)$. This is given by [19, Thm. 8.5]. Because $V(\xi) = \Gamma$ a.s., we have

$$K\big(\Gamma, \{\tau \in \mathcal{A}(S^2) : V(\tau) \neq \Gamma\}\big) = 0 \quad \text{a.s.}$$

Hence, this support condition fails only on a Γ-null set of γ's. Define

$$\Delta^2(\gamma) := \sum_{x \in \gamma} \delta_{(x,x)} \in \mathcal{A}(S^2),$$

and redefine K on that null set to deterministically be $\delta_{\Delta^2(\gamma)}$. Call the modified kernel C, which now satisfies

$$C(\gamma, \{\tau : V(\tau) \neq \gamma\}) = 0 \quad \forall \gamma \in \mathcal{N}(S).$$

The law of $C(\Gamma, \cdot)$ is unchanged, so $\xi \sim C(\Gamma, \cdot)$. ◀

This is evidence that the connection kernel perspective is not something distinct from our approach, but rather a means through which we can apply the tools that we have established. It allows one to specify a constrained and usable class of connection kernels, and verify the properties necessary for our theorems on such a class, which is exactly what we will soon do with spatial inhomogeneous random graphs.

However, before we proceed to our models, we first want to cover some general results for stationarity on connection kernels. This is because we would like the Palm transformation of the vertex process to construct the Palm transformation of the adjacency process under the same connection kernel. This allows us to shift model complexity into the connection kernel and avoid having to directly derive complex Palm distributions from the adjacency processes themselves. The necessary condition for this is known as equivariance of the connection kernel under the group action. Specifically, when considering vertex processes on a grouped LC Polish space (S, H), we say that a connection kernel C is *H-equivariant* if for every $\gamma \in \mathcal{N}(S)$ and $h \in H$, it holds that

$$C(\theta_h \gamma, \cdot) = \theta_h C(\gamma, \cdot), \tag{24}$$

where the action on S^2 is given by $h \cdot (x, y) \mapsto (hx, hy)$, just as in Section 5. The notation

$$\theta_h C(\gamma, \cdot)$$

is defined per $A \in \mathcal{B}(\mathcal{N}(S^2))$ as

$$\theta_h C(\gamma, A) := C(\gamma, \theta_h^{-1} A).$$

Essentially this condition guarantees that shifting the vertex process and then constructing the adjacency process is equivalent to shifting the original adjacency process. When we have this condition satisfied, we can associate stationarity and Palm transformations between the vertex and adjacency processes, as seen below.

▶ **Proposition 6.2.** *Let $(S, H)^2$ be a square grouped LC Polish space, with H acting transitively on S, and fix some $x \in S$. Let C be an H-equivariant connection kernel and let Γ be a vertex process on S. Consider the adjacency process $\xi \sim C(\Gamma, \cdot)$. Then,*

1. *Γ is stationary on (S, H) if and only if ξ is stationary on $(S, H)^2$.*
2. *In the case of such stationarity, $\xi^{(x,x)} \sim C(\Gamma^x, \cdot)$.*

It should be noted that it is possible for both ξ and Γ to be stationary without C being H-equivariant. However, since proper Palm behaviour is essential for us, we need equivariance as well, and it is a cleaner convention to use in our results.

6.2 Spatial Inhomogeneous Random Graphs

With the notion of a connection kernel complete, we can finally give a definition for the class of spatial inhomogeneous random graphs, denoted shortly as SIRGs. This will be done by defining a special type of connection kernel, called a *marked-isometric* connection kernel. Any random graph construction using such a kernel will be called a SIRG.

A marked-isometric connection kernel is a connection kernel that is broken down into two parts. First, it assigns the random collection of vertices each a random mark (measurably, according to a kernel), and then independently connects pairs with probabilities determined only by their distance and their marks. See the formal statement below.

Definition 22

Let S^2 be a square LC Polish space, and let $\mathbb{M}$ be an LC Polish space. Let $\Gamma \in \mathcal{N}(S)$ be a vertex process on S. Let $C : \mathcal{N}(S) \times \mathcal{B}(\mathcal{A}(S^2)) \to [0,1]$ be a connection kernel. We say that this connection kernel is *marked-isometric* iff there exist

1. a Markov kernel $M : \mathcal{N}(S) \times \mathcal{B}(\mathcal{N}(S \times \mathbb{M})) \to [0,1]$,
2. a measurable function $\kappa : \mathbb{R}^+ \times \mathbb{M} \times \mathbb{M} \to [0,1]$,
3. and a Markov kernel $Q_\kappa : \mathcal{N}(S \times \mathbb{M}) \times \mathcal{B}(\mathcal{A}(S^2)) \to [0,1]$,

satisfying the following properties:

1. $\kappa(t, m, m') = \kappa(t, m', m)$ for all $t \in \mathbb{R}^+$ and $m, m' \in \mathbb{M}$.
2. $\kappa(0, \cdot, \cdot) \equiv 0$.
3. For every $\gamma \in \mathcal{N}(S)$, it holds that $M(\gamma, A) = 0$, where
$$A = \left\{ \mu \in \mathcal{N}(S \times \mathbb{M}) \;\middle|\; \gamma \neq \sum_{(x,w_x)\in\mu} \delta_x \right\}.$$
4. For any $\mu \in \mathcal{N}(S \times \mathbb{M})$, and $\xi_\mu \sim Q_\kappa(\mu, \cdot)$, it holds that
$$\xi_\mu \stackrel{d}{=} \sum_{(x,m_x)\in\mu} \delta_{(x,x)} + \sum_{(x,m_x),(y,m_y)\in\mu}^{x\neq y} E_{\{x,y\}}\delta_{(x,y)},$$
where, for each distinct unordered pair of points $(x, m_x), (y, m_y) \in \mu$, the variable $E_{\{x,y\}}$ is an independent Bernoulli variable with probability $\kappa(d_S(x,y), m_x, m_y)$.
5. For every $\gamma \in \mathcal{N}(S)$,
$$C(\gamma, \cdot) = \int_{\mathcal{N}(S\times\mathbb{M})} Q_\kappa(\mu, \cdot) M(\gamma, d\mu).$$

If C is a marked-isometric connection kernel, then for any vertex process Γ on S, we call the adjacency process $\xi \sim C(\Gamma, \cdot)$ a *spatial inhomogeneous random graph*, and use the shorthand $\xi \sim \mathrm{SIRG}(\Gamma, M, \kappa)$ to identify the components.

We refer to M as a *marking kernel*, and κ as a *connection function*. Q_κ is given no special name.

It can be seen that we again have vertex support on M (by assumption) and on Q (by construction) to avoid changes to the vertex collection. Thus, the kernel C is indeed a connection kernel. This follows directly from composition of Markov kernels that both satisfy the same support condition (see [34, Def. 31] for composition of Markov kernels).

This is a very broad construction, as the marking kernel is where the true randomness occurs. In this regard, a marked-isometric connection kernel is not any more usable than a connection kernel when taken at full generality. However, if we ensure that a marking kernel M on a grouped LC Polish space (S, H) is H-equivariant, the model class becomes far more workable. Formally, this means that for every $h \in H$, $\gamma \in \mathcal{N}(S)$, it holds that

$$M(\theta_h \gamma, \cdot) = \theta_h M(\gamma, \cdot) \tag{25}$$

where the action of H on $S \times \mathbb{M}$ is given by $(x, m) \mapsto (hx, m)$. This is analogous to the definition of H-equivariance for connection kernels, and in fact, it implies equivariance of the connection kernels whenever the group action is isometric. Isometric here means that the group action preserves distance, i.e. for any $h \in H$ and $x, y \in S$, it holds that

$$d_S(hx, hy) = d_S(x, y).$$

See this lemma below.

▶ **Lemma 6.3.** *Let $(S, H)^2$ be a square grouped LC Polish space and $\mathbb{M}$ be an LC Polish space. Let $\xi \sim \mathrm{SIRG}(\Gamma, M, \kappa)$ be a SIRG on S^2 taking marks in $\mathbb{M}$. Suppose that H acts transitively and isometrically on S. Then, C as in Definition 22 is H-equivariant whenever M is H-equivariant.*

The proof is again given in Section 6.5 and allows us to transfer the stationarity and Palm results from Proposition 6.2 to SIRGs. For simplicity, we will refer to a SIRG with H-equivariant marking kernel as simply an *H-equivariant SIRG*, and if the group action is clear from context, we simply say *equivariant SIRG*. When we say stationary SIRG, we will mean that both its vertex and adjacency processes are stationary. Of course, we still need an isometric action as well to apply the result. On this note, we move to defining the sequences of spaces that will be admissible for our later results.

6.3 β-Admissible Sequences of Spaces

Just as in Section 5, we need to be cautious when working with sequences. For our use case, we will need three additional conditions on top of all of the α-admissibility conditions. The first is isometry of all the group actions, which is simple. The other two are restrictions on the behaviour of the metrics, that ensure that the sequence of spaces eventually has metrics that behave "similarly" to the ambient space. These conditions are essential because SIRGs fundamentally work with distance. We give the definition, where again β is just a disambiguator.

Definition 23 β-Admissible Sequence of Spaces

Let S^2 be a square LC Polish space, and fix some $x \in S$. Let

$$(S_n, H_n)^2_{n\in\mathbb{N}} \subset_{\alpha,x} S^2$$

be an α-admissible sequence of spaces. Suppose the following conditions also hold:

1. H_n acts isometrically on S_n for every $n \in \mathbb{N}$;
2. The metrics d_{S_n} converge sequentially to d_S, i.e. for all sequences $(x_n)_{n\in\mathbb{N}}, (y_n)_{n\in\mathbb{N}}$ with $x_n, y_n \in S_n$ converging to $x, y \in S$, it holds that
$$d_{S_n}(x_n, y_n) \to d_S(x, y);$$
3. For all $R > 0$, there exists a compact $K_R \subset S$ with $B_{d_{S_n}}(x, R) \subset K_R$ eventually.

Then, we call $(S_n, H_n)^2_{n\in\mathbb{N}}$ an *β-admissible sequence of spaces* in S^2 sharing $(x, x) \in S^2$, and denote it as $(S_n, H_n)^2_{n\in\mathbb{N}} \subset_{\beta,x} S^2$.

Condition (2) is a metric-compatibility requirement: distances computed in the approximating spaces should agree asymptotically with distances in the ambient space whenever the underlying points converge. This is needed so that distance-based connection rules in SIRGs have a stable limit interpretation. Condition (3) is an eventual compact-containment requirement for metric balls around the base point: fixed-radius neighborhoods in the approximating spaces must not drift unboundedly in the ambient space. In the tightness proof, this is exactly what lets us replace $B_{d_{S_n}}(x, R)$ by one fixed compact set in S, uniformly for all large n.

The notion of sequential convergence from condition (2) is something that we will use again further down, on our connection functions. We say that a sequence of functions $(f_n)_{n\in\mathbb{N}}$ on a shared domain D converges *sequentially* to a function f if for every sequence $(x_n)_{n\in\mathbb{N}}$ in the D converging to some $x \in D$, it holds that

$$f_n(x_n) \to f(x)$$

whenever x is a continuity point of f. We denote this as

$$f_n \xrightarrow{s} f.$$

The fact that this is only required for continuity points is noteworthy here, as it improves generality on connection functions. For condition (2) in the β-admissibility definition, this is redundant to state, as metrics are always continuous with respect to their own metric space.

6.4 Local Weak Convergence Theorem for SIRGs

In order to state our local convergence theorem, we need a few definitions. The first is an analog of sequential convergence, tailored to kernels. Let $(K_n)_{n\in\mathbb{N}}, K : X \times \mathcal{B}(Y) \to [0,1]$ be a collection of kernels. We say that the sequence $(K_n)_{n\in\mathbb{N}}$ converges sequentially weakly to K if for every sequence x_n in X converging to some $x \in X$, it holds that

$$K_n(x_n) \Rightarrow K(x)$$

whenever x is a continuity point of K. Here, the $\Rightarrow$ denotes weak convergence of laws, i.e. convergence against all bounded continuous test functions. We denote sequential weak convergence as

$$K_n \overset{s}{\Longrightarrow} K.$$

The second necessary definition is that of the second factorial moment measure. This is essentially the intensity measure of ordered pairs of distinct points of a point process. In other words, it records the expected number of pairs of different points falling in two prescribed sets.

Formally, if η is a point process on an LC Polish space S, its second factorial moment measure is the measure $\alpha_\eta^{(2)}$ on S^2 defined for measurable $A, B \subseteq S$ as

$$\alpha_\eta^{(2)}(A \times B) = \mathbb{E}\left[\sum_{\substack{u,v\in\eta \\ u\neq v}} \mathbf{1}_A(u)\mathbf{1}_B(v)\right]. \tag{26}$$

Thus, $\alpha_\eta^{(2)}(A \times B)$ is the expected number of pairs (u, v) with $u \in A$ and $v \in B$. Notably, it counts both orderings (u, v) and (v, u). This measure is necessary, as it helps us to understand the behavior of jointly occurring points, which is essential for tightness and fiber-finiteness criteria. The standard intensity measure does not provide enough information to arrive at these results. The second factorial moment also has a useful functional identity. For every non-negative measurable function $f : S \times S \to \mathbb{R}^+$,

$$\mathbb{E}\left[\sum_{x,y\in\eta}^{x\neq y} f(x,y)\right] = \iint_{S^2} f(x,y)\ \alpha_\eta^{(2)}(dx,dy), \tag{27}$$

which we will use in our later proofs. We also need to establish the notion of an identically factorable marking kernel. We call a marking kernel M *identically factorable* with law ν if for any $\gamma \in \mathcal{N}(S)$ and $\mathcal{X} \sim M(\gamma, \cdot)$, it holds that

$$\mathcal{X} \overset{d}{=} \sum_{x\in\gamma} \delta_{(x,W_x)},$$

where the family $(W_x)_{x\in\gamma}$ is an i.i.d. family with law ν on the mark space. When we have a connection kernel κ, we can then define the mark-averaged

connection kernel,

$$\overline{\kappa}^{\nu}(t) := \iint_{\mathbb{M}\times\mathbb{M}} \kappa(t, m, m')\ \nu(dm)\ \nu(dm').$$

When the law is clear from context, we simply use the notation $\overline{\kappa}$. This is essential because we can reduce radial tightness and fiber-finiteness checks to this averaged connection kernel, rather than the full law. While it is likely possible that only the limiting kernel needs to be identically factorable, the scope of this paper relies on every marking kernel in the sequence being identically factorable to achieve the central results.

The last notion we need to establish before the theorem is that of the distance process. This is because we need to ensure that the distances between realized points do have positive probability singleton values at discontinuity points of the limiting connection kernel. For a point process η on an LC Polish space S, the corresponding distance process $D(\eta)$ is the point process

$$D(\eta) := \sum_{x,y\in\eta}^{x\neq y} \delta_{d_S(x,y)}.$$

While this is not a simple point process, multiplicities are irrelevant because we are only using it to check continuity points. We only count distinct points because self distances are irrelevant for our continuity checks. With all of these defined, we can give the local weak convergence theorem for SIRGs.

▶ **Theorem 6.4** Local Weak Convergence Theorem for SIRGs. *Let S^2 be a square LC Polish space and $\mathbb{M}$ be an LC Polish space, and fix $x \in S$. Let*

$$(\xi_n)_{n\in\mathbb{N}} \sim \mathrm{SIRG}(\Gamma_n, M, \kappa_n)_{n\in\mathbb{N}}$$

be a sequence of finite stationary equivariant SIRGs on a β-admissible sequence of spaces $(S_n, H_n)^2 \subset_{\beta,x} S^2$ taking marks in $\mathbb{M}$. Suppose that M is an identically factorable marking kernel with law ν, and let W_1, W_2 be i.i.d. random variables with distribution ν. Let also

$$\xi_\infty \sim \mathrm{SIRG}(\Gamma_\infty, M, \kappa_\infty)$$

be a SIRG for some vertex process Γ_∞ and some connection function κ_∞. Suppose that all of the following hold:

1. *$\Gamma_n^x \Rightarrow \Gamma_\infty$ in $\mathcal{N}(S)$ and $\kappa_n \to \kappa_\infty$ sequentially.*
2. *κ_∞ is $(D(\Gamma_\infty)\otimes\nu\otimes\nu)$-a.s. continuous.*
3. *The vertex counts concentrate via*

$$\mathbb{E}\left[\left|\frac{|\Gamma_n|}{\mathbb{E}[|\Gamma_n|]} - 1\right|\right] \to 0.$$

4. *For every compact set $L \subseteq S$,*

$$\int_{L\times S} \overline{\kappa}(d_S(u,v))\ \alpha^{(2)}_{\Gamma_\infty}(du, dv) < \infty,$$

5. *For all compact $L \subseteq S$, there exists some $N \in \mathbb{N}$ for which*

$$\lim_{R\to\infty} \sup_{n\geq N} \iint_{L\times(S\setminus B_S(x,R))} \overline{\kappa}_n(d_{S_n}(u,v))\ \alpha^{(2)}_{\Gamma_n^x}(du,dv) = 0.$$

Then, it also holds that

$$G(\xi_n, U_n) \Rightarrow G(\xi_\infty, x),$$

where $U_n \sim \mathbb{U}(\xi_n, \cdot)$.

This theorem is a structured checklist for Theorem 5.2. Conditions (1) and (2) give convergence of $(\xi_n^{(x,x)})_{n\in\mathbb{N}}$ to ξ_∞, condition (3) gives vertex concentration, condition (4) gives fiber-finiteness, and condition (5) gives radial tightness. We now state the intermediate results that supply each of these items. The first is our continuity result.

▶ **Proposition 6.5.** *Let S^2 be a square LC Polish space and $\mathbb{M}$ be an LC Polish space, and fix $x \in S$. Let*

$$(\xi_n)_{n\in\mathbb{N}} \sim \mathrm{SIRG}(\Gamma_n, M_n, \kappa_n)_{n\in\mathbb{N}}$$

be a sequence of SIRGs on a β-admissible sequence of spaces $(S_n, H_n)^2 \subset_{\beta,x} S^2$ taking marks in $\mathbb{M}$. Suppose that $M_n \Rightarrow^s M_\infty$ sequentially weakly for some limit identically factorable marking kernel M_∞ with law ν. Suppose that $\kappa_n \to \kappa_\infty$ sequentially to some limit connection function κ_∞, and that κ_∞ is $(D(\Gamma_\infty) \otimes \nu \otimes \nu)$-a.s. continuous. Then, whenever $\Gamma_n \Rightarrow \Gamma_\infty$, it holds that

$$\xi_n \Rightarrow \xi_\infty$$

for $\xi_\infty \sim \mathrm{SIRG}(\Gamma_\infty, M_\infty, \kappa_\infty)$.

This is a general continuity result for marked-isometric connection kernels, and thus for SIRGs. Under these conditions, determining convergence of the vertex process and relevant components is sufficient to determine convergence of the adjacency processes and construct the limit. This also shows why we needed the sequence metric convergence for β-admissibility, as without it distances of corresponding pairs can drift, so the connection functions need not stabilize and edge probabilities can fail to converge. At this level, it is still possible to have sequential weak convergence of kernels instead of forcing all to be identically factorable. In fact, the only place where the kernels all need to be identically factorable is the current condition for radial tightness, which will be seen shortly. But first, we give a condition for fiber-finiteness, and thus regularity, of infinite SIRGs.

▶ **Lemma 6.6.** *Let S^2 be a square LC Polish space and $\mathbb{M}$ be an LC Polish space. Let $\xi \sim \mathrm{SIRG}(\Gamma, M, \kappa)$ be a SIRG taking marks in $\mathbb{M}$, and suppose that M is an identically factorable marking kernel with law ν. Then, if it holds that for every compact set $L \subseteq S$,*

$$\int_{L\times S} \overline{\kappa}(d_S(x,y))\ \alpha^{(2)}_{\Gamma}(dx,dy) < \infty,$$

it holds that ξ is a fiber-finite adjacency process.

This result will be primarily used to guarantee that our limiting SIRG is regular, since our sequences are of finite SIRGs, and thus do not require it. Essentially, by Equation (27), this integral is exactly the expected number of incident edge-pairs touching a compact set, so its finiteness enforces local edge finiteness and hence fiber-finiteness. The final condition that we need is our criterion for radial tightness. See the statement below.

▶ **Lemma 6.7.** *Let S^2 be a square LC Polish space and $\mathbb{M}$ be an LC Polish space, and fix $x \in S$. Let*

$$(\xi_n)_{n\in\mathbb{N}} \sim \mathrm{SIRG}(\Gamma_n, M, \kappa_n)_{n\in\mathbb{N}}$$

be a sequence of finite stationary equivariant SIRGs on a β-admissible sequence of spaces $(S_n, H_n)^2 \subset_{\beta,x} S^2$ taking marks in $\mathbb{M}$. Suppose that M is identically factorable with law ν. Then, if for all compact $L \subseteq S$ there exists some $N \in \mathbb{N}$ for which

$$\lim_{R\to\infty} \sup_{n\geq N} \iint_{L\times(S\setminus B_S(x,R))} \overline{\kappa}_n(d_{S_n}(u,v))\, \alpha^{(2)}_{\Gamma_n^x}(du, dv) = 0,$$

it holds that $(\xi_n^{(x,x)}, x)$ is a radially tight sequence.

Broadly, this integral measures the expected number of Palm-rooted edges leaving a fixed compact core L and reaching beyond distance R. This becomes uniformly negligible for all large n, which prevents mass from escaping far from the root and is exactly what the proof converts into radial tightness of rooted neighborhoods.

As can be seen here, this result requires all marking kernels to be identically factorable with the same law, so that we can use the same mark-averaged connection function in the supremum tail condition. However, this result is the only bottleneck that forces this, and it is likely that it can be generalized in future work. With these done, we give the proof of Theorem 6.4 up to these intermediate results.

Proof of Theorem 6.4

Proof. We apply Theorem 5.2 from Section 5 to the sequence $(\xi_n)_{n\in\mathbb{N}}$. Since

$$(S_n, H_n)^2_{n\in\mathbb{N}} \subset_{\beta,x} S^2,$$

it is in particular α-admissible by Definition 23. Thus, it remains to verify the three hypotheses of Theorem 5.2.

First, we verify convergence of Palm variants to a regular limit. Because each ξ_n is a stationary equivariant SIRG, Lemma 6.3 and Proposition 6.2 gives

$$\xi_n^{(x,x)} \sim \mathrm{SIRG}(\Gamma_n^x, M, \kappa_n).$$

Now apply Proposition 6.5 to this Palm-transformed sequence: assumption (1) gives $\Gamma_n^x \Rightarrow \Gamma_\infty$ and sequential convergence $\kappa_n \to \kappa_\infty$, while assumption (2) gives the required $(D(\Gamma_\infty) \otimes \nu \otimes \nu)$-a.s. continuity of κ_∞. Since the marking kernel is fixed, we take $M_n \equiv M$, so $M_n \overset{s}{\Longrightarrow} M$ trivially. Therefore,

$$\xi_n^{(x,x)} \Rightarrow \xi_\infty,$$

where $\xi_\infty \sim \mathrm{SIRG}(\Gamma_\infty, M, \kappa_\infty)$. Also, since each Palm-rooted vertex process Γ_n^x has a point at x almost surely, the support condition in Definition 22 gives that each $\xi_n^{(x,x)}$ has the diagonal atom (x, x) almost surely, and hence the limit ξ_∞ is non-empty. By assumption (4) and Lemma 6.6, ξ_∞ is fiber-finite, hence regular. Additionally, assumption (5) is exactly the hypothesis of Lemma 6.7, so

$$(\xi_n^{(x,x)}, x)_{n\in\mathbb{N}}$$

is radially tight. Lastly, the vertex condition in Theorem 5.2 follows from assumption (3): by the support condition in Definition 22, $V(\xi_n) = \Gamma_n$ almost surely for each n, so

$$\mathbb{E}\left[\left|\frac{|V(\xi_n)|}{\mathbb{E}[|V(\xi_n)|]} - 1\right|\right] = \mathbb{E}\left[\left|\frac{|\Gamma_n|}{\mathbb{E}[|\Gamma_n|]} - 1\right|\right] \to 0.$$

All hypotheses of Theorem 5.2 are now verified, hence

$$G(\xi_n, U_n) \Rightarrow G(\xi_\infty, x),$$

where $U_n \sim \mathbb{U}(\xi_n, \cdot)$. ◀

6.5 Propositions

We start with the proof of Proposition 6.2, as it is the first to appear linearly. The argument is direct. We first test stationarity against a bounded measurable functional, use equivariance to move the group action from the adjacency process back to the vertex process, and then apply stationarity of the vertex process. The reverse direction follows by projecting the adjacency process back to its vertex set. The Palm statement is then obtained by counting diagonal atoms in two equivalent ways and then using uniqueness of the resulting Palm disintegration.

Proof of Proposition 6.2

Proof. We first show that stationarity is equivalent. First assume that Γ is stationary. Let $F : \mathcal{A}(S^2) \to \mathbb{R}$ be bounded measurable and $h \in H$. Using $\xi \sim C(\Gamma, \cdot)$,

$$\mathbb{E}[F(\theta_h \xi)] = \mathbb{E}\left[\int_{\mathcal{A}(S^2)} F(\theta_h \tau)\, C(\Gamma, d\tau)\right].$$

By equivariance of C,

$$\int_{\mathcal{A}(S^2)} F(\theta_h\tau)\, C(\gamma, d\tau) = \int_{\mathcal{A}(S^2)} F(\tau)\, C(\theta_h\gamma, d\tau).$$

Hence,

$$\mathbb{E}[F(\theta_h\xi)] = \mathbb{E}\left[\int F(\tau)\, C(\theta_h\Gamma, d\tau)\right].$$

Since Γ is stationary by assumption, $\theta_h\Gamma \stackrel{d}{=} \Gamma$, so

$$\mathbb{E}[F(\theta_h\xi)] = \mathbb{E}\left[\int F(\tau)\, C(\Gamma, d\tau)\right] = \mathbb{E}[F(\xi)].$$

Thus, ξ is stationary on $(S,H)^2$. Alternatively, suppose that ξ is stationary. Then, since C is a connection kernel, $V(\xi) = \Gamma$ a.s. Also, $V(\theta_h\tau) = \theta_h V(\tau)$. Therefore,

$$\theta_h\Gamma = \theta_h V(\xi) = V(\theta_h\xi) \stackrel{d}{=} V(\xi) = \Gamma,$$

and Γ is stationary on (S,H). It remains only to verify the Palm relation. Assume shared stationarity of the processes. The Palm versions at x and (x,x) exist due to the assumed transitivity. Take nonnegative measurable $g : S \times \mathcal{A}(S^2) \to [0,\infty)$. First, using $\xi \sim C(\Gamma,\cdot)$, we have that

$$\mathbb{E}\left[\sum_{u\in\Gamma} g(u,\xi)\right] = \mathbb{E}\left[\sum_{u\in\Gamma}\int g(u,\tau)\, C(\Gamma, d\tau)\right].$$

Applying the Palm identity for Γ then gives

$$\mathbb{E}\left[\sum_{u\in\Gamma} g(u,\xi)\right] = \int_S \mathbb{E}^x_\Gamma\left[\int g(x,\tau)\, C(\Gamma, d\tau)\right]\Lambda_\Gamma(dx). \tag{28}$$

On the other hand, note that by $V(\xi) = \Gamma$ almost surely it holds that

$$\sum_{u\in\Gamma} g(u,\xi) = \sum_{z\in\xi} \mathbf{1}_{\{z\in\Delta^2(S)\}}\, g(\pi_1(z),\xi).$$

Applying the Palm identity for ξ with

$$f(z,\tau) := \mathbf{1}_{\{z\in\Delta^2(S)\}}\, g(\pi_1(z),\tau)$$

gives

$$\mathbb{E}\left[\sum_{u\in\Gamma} g(u,\xi)\right] = \int_{S^2} \mathbf{1}_{\{z\in\Delta^2(S)\}}\, \mathbb{E}^z_\xi[g(\pi_1(z),\xi)]\, \Lambda_\xi(dz).$$

Now for any Borel set $A \subseteq S$,

$$\Lambda_\xi(\Delta^2(A)) = \mathbb{E}[\xi(\Delta^2(A))] = \mathbb{E}[\Gamma(A)] = \Lambda_\Gamma(A),$$

since the diagonal atoms of ξ are exactly the vertices of Γ almost surely. Therefore, the restriction of Λ_ξ to $\Delta^2(S)$ is the pushforward of Λ_Γ under $x \mapsto (x, x)$, and so

$$\mathbb{E}\left[\sum_{u\in\Gamma} g(u,\xi)\right] = \int_S \mathbb{E}_\xi^{(x,x)}[g(x,\xi)]\,\Lambda_\Gamma(dx). \tag{29}$$

To pass from these two integral identities to a pointwise statement at the chosen basepoint, define

$$F(y) := \mathbb{E}_\xi^{(y,y)}[g(y,\xi)], \qquad G(y) := \mathbb{E}_\Gamma^{y}\left[\int g(y,\tau)\,C(\Gamma,d\tau)\right].$$

If we replace $g(u,\tau)$ in the preceding argument by $\mathbf{1}_A(u)\,g(u,\tau)$ for an arbitrary Borel set $A \subseteq S$, then Equation (28) and Equation (29) give

$$\int_A F(y)\,\Lambda_\Gamma(dy) = \int_A G(y)\,\Lambda_\Gamma(dy).$$

Since this holds for every Borel A, it follows that $F = G$ for Λ_Γ-almost every $y \in S$. Now H acts transitively on S, and stationarity of Γ and ξ together with Palm covariance along the diagonal orbit imply that both Palm kernels are determined on the whole orbit $Hx = S$ from their values on any Λ_Γ-non-null subset. Hence, this almost-everywhere identity extends to the distinguished basepoint x, giving

$$\mathbb{E}_\xi^{(x,x)}[g(x,\xi)] = \mathbb{E}_\Gamma^{x}\left[\int g(x,\tau)\,C(\Gamma,d\tau)\right].$$

for every nonnegative measurable $g : S\times\mathcal{A}(S^2) \to [0,\infty)$. Taking $g(u,\tau) = f(\tau)$, where $f : \mathcal{A}(S^2) \to [0,\infty)$ is measurable, gives

$$\mathbb{E}_\xi^{(x,x)}[f(\xi)] = \mathbb{E}_\Gamma^{x}\left[\int f(\tau)\,C(\Gamma,d\tau)\right].$$

Now define probability measures on $\mathcal{A}(S^2)$ by

$$\mu_x(B) := \mathbb{P}_\xi^{(x,x)}(\xi \in B), \qquad \nu_x(B) := \mathbb{E}_\Gamma^{x}[C(\Gamma,B)], \qquad B \in \mathcal{B}(\mathcal{A}(S^2)).$$

The displayed identity is exactly

$$\int f\,d\mu_x = \int f\,d\nu_x$$

for all nonnegative measurable f. Choosing $f = \mathbf{1}_B$ yields $\mu_x(B) = \nu_x(B)$ for every Borel set B, hence $\mu_x = \nu_x$. Therefore,

$$\xi^{(x,x)} \sim C(\Gamma^x,\cdot).$$

Which concludes the proof. ◀

The other proposition that needs proof is Proposition 6.5. At a high level, the argument has two layers. First, we fix an arbitrary convergent deterministic sequence of vertex configurations and show that the corresponding adjacency laws converge. To do this, we compare Laplace functionals, reduce to the finitely many vertices that can influence a compactly supported test function, and pass to the limit in the explicit independent edge-factorization, using continuity of the limiting connection rule at the relevant realized distance and mark triples. This gives sequential convergence of the composed connection kernels for every deterministic convergent input sequence. Second, we apply the varying-map step to the random vertex processes, which upgrades the deterministic sequential statement to convergence in distribution of the full adjacency processes.

Proof of Proposition 6.5

Proof. First note that κ_∞ has the defining structural identities of a connection function from Definition 22. For each $n \in \mathbb{N}$, write $Q_n := Q_{\kappa_n}$ and define

$$C_n(\gamma, \cdot) := \int_{\mathcal{N}(S\times\mathbb{M})} Q_n(\mu, \cdot)\, M_n(\gamma, d\mu),$$

and analogously

$$C_\infty(\gamma, \cdot) := \int_{\mathcal{N}(S\times\mathbb{M})} Q_\infty(\mu, \cdot)\, M_\infty(\gamma, d\mu).$$

Since composition of Markov kernels is again a Markov kernel, and both M_n and Q_n satisfy the support condition from Definition 22, each C_n (including C_∞) is a connection kernel. We use Laplace functionals throughout. Write

$$\mathcal{C}_c^+(S) := \{h : S \to [0,\infty) \,:\, h \text{ is continuous with compact support}\},$$

and analogously $\mathcal{C}_c^+(S \times \mathbb{M})$. For a point process η on S and $h \in \mathcal{C}_c^+(S)$, set $L_\eta(h) := \mathbb{E}[e^{-\langle h,\eta\rangle}]$, where here and below,

$$\langle h, \eta\rangle := \int h\, d\eta,$$

over the ambient space of η. Under the vague topology, convergence against these test functionals is sufficient to determine weak convergence by [2, Sec. 4, p. 5].

We first note that M_∞ is continuous at every $\gamma \in \mathcal{N}(S)$. Indeed, let $f \in \mathcal{C}_c^+(S \times \mathbb{M})$, set

$$a_f(x) := -\log\!\left(\int_{\mathbb{M}} e^{-f(x,m)}\, \nu(dm)\right), \qquad x \in S.$$

Then $a_f \in \mathcal{C}_c^+(S)$, and for $\mathcal{X}_\gamma \sim M_\infty(\gamma, \cdot)$ (identically factorable with law ν),

$$\mathbb{E}\Big[e^{-\langle f,\mathcal{X}_\gamma\rangle}\Big] = \prod_{x\in\gamma} \int_{\mathbb{M}} e^{-f(x,m)}\, \nu(dm) = \exp\big(-\langle a_f, \gamma\rangle\big),$$

where the product is finite away from 1 because a_f has compact support. Hence, $\gamma_n \to \gamma$ implies $\langle a_f, \gamma_n \rangle \to \langle a_f, \gamma \rangle$, so the Laplace functionals converge for every $f \in \mathcal{C}_c^+(S \times \mathbb{M})$, and therefore

$$M_\infty(\gamma_n, \cdot) \Rightarrow M_\infty(\gamma, \cdot).$$

So every γ is a continuity point of M_∞. We now verify sequential weak convergence $C_n \overset{s}{\Longrightarrow} C_\infty$. Fix any deterministic sequence

$$\gamma_n \to \gamma_\infty \quad \text{in } \mathcal{N}(S),$$

and let

$$\mathcal{X}_n \sim M_n(\gamma_n, \cdot), \qquad \mathcal{X}_\infty \sim M_\infty(\gamma_\infty, \cdot).$$

By the assumed sequential weak convergence $M_n \overset{s}{\Longrightarrow} M_\infty$, together with continuity of M_∞ at every point, we have $\mathcal{X}_n \Rightarrow \mathcal{X}_\infty$. Fix $g \in \mathcal{C}_c^+(S^2)$, and define

$$\psi_{n,g}(\mu) := \int_{\mathcal{N}(S^2)} e^{-\langle g, \eta \rangle}\, Q_n(\mu, d\eta), \qquad \psi_{\infty,g}(\mu) := \int_{\mathcal{N}(S^2)} e^{-\langle g, \eta \rangle}\, Q_\infty(\mu, d\eta).$$

These maps are measurable and $[0,1]$-valued, since $g \geq 0$ implies $0 \leq e^{-\langle g, \eta \rangle} \leq 1$ for every $\eta \in \mathcal{N}(S^2)$, the map $\eta \mapsto e^{-\langle g, \eta \rangle}$ is measurable, and integrating a bounded measurable function against the probability kernels $Q_n(\mu, \cdot)$ and $Q_\infty(\mu, \cdot)$ is measurable in μ. By construction,

$$\mathbb{E}[\psi_{n,g}(\mathcal{X}_n)] = \int_{\mathcal{N}(S^2)} e^{-\langle g, \eta \rangle}\, C_n(\gamma_n, d\eta),$$

$$\mathbb{E}[\psi_{\infty,g}(\mathcal{X}_\infty)] = \int_{\mathcal{N}(S^2)} e^{-\langle g, \eta \rangle}\, C_\infty(\gamma_\infty, d\eta).$$

Let

$$L := \pi_1(\operatorname{supp} g) \cup \pi_2(\operatorname{supp} g) \subset S,$$

which is compact. Define the continuity set of κ_∞,

$$\mathcal{C}_\infty := \{(r, m, m') : \ \kappa_\infty \text{ is continuous at } (r, m, m')\},$$

and the good marked configurations

$$\mathcal{G}_g := \Big\{ \mu \in \mathcal{N}(S \times \mathbb{M}) \;\Big|\; \text{for all } (x, m_x), (y, m_y) \in \mu \text{ with } x, y \in L,\ x \neq y, \\ (d_S(x, y), m_x, m_y) \in \mathcal{C}_\infty \Big\}.$$

Take any deterministic sequence $\mu_n \to \mu_\infty$ in $\mathcal{N}(S \times \mathbb{M})$ with $\mu_\infty \in \mathcal{G}_g$. Only atoms with spatial coordinate in L contribute to $\langle g, \xi_\mu \rangle$, and there are finitely many such atoms since μ is boundedly finite and L is compact. By the standard

atom-matching property, [2, Prop. 2.8], for simple counting measures, for all large n we can label these atoms as

$$\{(x_n^1, m_n^1), \dots, (x_n^J, m_n^J)\} \subset \mu_n, \qquad \{(x^1, m^1), \dots, (x^J, m^J)\} \subset \mu_\infty,$$

with $(x_n^j, m_n^j) \to (x^j, m^j)$ for each j. By Definition 22, we can write

$$\xi_{\mu_n} = \sum_{j=1}^{J} \delta_{(x_n^j, x_n^j)} + \sum_{1 \le j < k \le J} E_n^{\{j,k\}} \big(\delta_{(x_n^j, x_n^k)} + \delta_{(x_n^k, x_n^j)}\big),$$

where $\{E_n^{\{j,k\}} : 1 \le j < k \le J\}$ are independent Bernoulli variables with

$$p_n^{jk} := \mathbb{P}(E_n^{\{j,k\}} = 1) = \kappa_n\big(d_{S_n}(x_n^j, x_n^k), m_n^j, m_n^k\big).$$

Hence, by the definition of integration and expectation rules,

$$\begin{aligned}
\psi_{n,g}(\mu_n) &= \mathbb{E}\Big[e^{-\langle g, \xi_{\mu_n}\rangle}\Big] \\
&= \mathbb{E}\left[\exp\left(-\sum_{j=1}^{J} g(x_n^j, x_n^j) - \sum_{1 \le j < k \le J} E_n^{\{j,k\}} \big(g(x_n^j, x_n^k) + g(x_n^k, x_n^j)\big)\right)\right] \\
&= \exp\left(-\sum_{j=1}^{J} g(x_n^j, x_n^j)\right) \prod_{1 \le j < k \le J} \mathbb{E}\Big[e^{-E_n^{\{j,k\}} \left(g(x_n^j, x_n^k) + g(x_n^k, x_n^j)\right)}\Big] \\
&= \exp\left(-\sum_{j=1}^{J} g(x_n^j, x_n^j)\right) \prod_{1 \le j < k \le J} \Big[1 - p_n^{jk}\Big(1 - e^{-\left(g(x_n^j, x_n^k) + g(x_n^k, x_n^j)\right)}\Big)\Big] \\
&= \exp\left(-\sum_{j=1}^{J} g(x_n^j, x_n^j)\right) \\
&\quad \times \prod_{1 \le j < k \le J} \Big[1 - \kappa_n\big(d_{S_n}(x_n^j, x_n^k), m_n^j, m_n^k\big)\Big(1 - e^{-\left(g(x_n^j, x_n^k) + g(x_n^k, x_n^j)\right)}\Big)\Big].
\end{aligned}$$

By β-admissibility, $d_{S_n}(x_n^j, x_n^k) \to d_S(x^j, x^k)$. Since $\mu_\infty \in \mathcal{G}_g$, each limiting triple $(d_S(x^j, x^k), m^j, m^k)$ is a continuity point of κ_∞, so sequential convergence $\kappa_n \to \kappa_\infty$ gives

$$\kappa_n\big(d_{S_n}(x_n^j, x_n^k), m_n^j, m_n^k\big) \to \kappa_\infty\big(d_S(x^j, x^k), m^j, m^k\big).$$

Also, $g(x_n^j, x_n^k) \to g(x^j, x^k)$ by continuity. Therefore, each factor converges, and since the product is finite,

$$\psi_{n,g}(\mu_n) \to \psi_{\infty,g}(\mu_\infty).$$

By identically factorable marking for M_∞ with law ν, together with the assumption that κ_∞ is $(D(\Gamma_\infty) \otimes \nu \otimes \nu)$-a.s. continuous, we obtain

$$\mathbb{P}(\mathcal{X}_\infty \in \mathcal{G}_g) = 1.$$

Applying the sequential mapping theorem [1, Lem. 1.3] for varying maps to $\mathcal{X}_n \Rightarrow \mathcal{X}_\infty$ and the convergence above on $\mathcal{G}_g$, we get

$$\psi_{n,g}(\mathcal{X}_n) \Rightarrow \psi_{\infty,g}(\mathcal{X}_\infty).$$

Since $0 \leq \psi_{n,g} \leq 1$, we may regard these as $[0,1]$-valued random variables; the identity map $t \mapsto t$ is bounded and continuous on $[0,1]$, so Proposition 2.1 gives

$$\mathbb{E}[\psi_{n,g}(\mathcal{X}_n)] \to \mathbb{E}[\psi_{\infty,g}(\mathcal{X}_\infty)].$$

Hence, for every $g \in \mathcal{C}_c^+(S^2)$,

$$\int_{\mathcal{N}(S^2)} e^{-\langle g,\eta\rangle}\, C_n(\gamma_n, d\eta) \longrightarrow \int_{\mathcal{N}(S^2)} e^{-\langle g,\eta\rangle}\, C_\infty(\gamma_\infty, d\eta),$$

and by the Laplace-functional criterion this implies

$$C_n(\gamma_n, \cdot) \Rightarrow C_\infty(\gamma_\infty, \cdot).$$

The convergent deterministic sequence $\gamma_n \to \gamma_\infty$ was arbitrary, and no extra condition on γ_∞ was imposed. Hence we have proved

$$C_n(\gamma_n, \cdot) \Rightarrow C_\infty(\gamma_\infty, \cdot) \quad \text{for every deterministic sequence } \gamma_n \to \gamma_\infty.$$

In particular, $C_n \overset{s}{\Longrightarrow} C_\infty$, and this statement applies to every possible realized value of the random limit Γ_∞. Let $F \in \mathcal{C}_b(\mathcal{N}(S^2))$, and set

$$h_n(\gamma) := \int_{\mathcal{N}(S^2)} F(\eta)\, C_n(\gamma, d\eta), \qquad h_\infty(\gamma) := \int_{\mathcal{N}(S^2)} F(\eta)\, C_\infty(\gamma, d\eta).$$

Therefore, for every deterministic $\gamma_n \to \gamma_\infty$,

$$h_n(\gamma_n) \to h_\infty(\gamma_\infty).$$

Applying the same varying-map theorem to $\Gamma_n \Rightarrow \Gamma_\infty$, we obtain

$$\mathbb{E}[F(\xi_n)] = \mathbb{E}[h_n(\Gamma_n)] \to \mathbb{E}[h_\infty(\Gamma_\infty)] = \mathbb{E}[F(\xi_\infty)].$$

Therefore, $\xi_n \Rightarrow \xi_\infty$, where

$$\xi_\infty \sim C_\infty(\Gamma_\infty, \cdot) = \mathrm{SIRG}(\Gamma_\infty, M_\infty, \kappa_\infty).$$

◀

6.6 Lemmas

The first lemma that we prove is Lemma 6.3. This is what allows us to apply the stationarity and Palm result, Proposition 6.2, to SIRGs. The proof is straightforward. We first check that shifting a marked vertex configuration before constructing the SIRG gives the same law as constructing the SIRG first and then shifting the resulting adjacency process. This follows because the group action preserves distances and leaves the marks unchanged, so every Bernoulli edge variable keeps the same parameter. Equivariance of the marking kernel then lets us pass this identity through via composition of kernels.

Proof of Lemma 6.3

Proof. Fix $h \in H$. Rewrite the two shift operations, for proof clarity, as

$$\theta_h^{(1)}\mu := \sum_{(x,m)\in\mu} \delta_{(hx,m)} \quad \text{on } \mathcal{N}(S\times\mathbb{M}), \qquad \theta_h^{(2)}\tau := \sum_{(x,y)\in\tau} \delta_{(hx,hy)} \quad \text{on } \mathcal{A}(S^2).$$

First we show

$$Q_\kappa(\theta_h^{(1)}\mu,\cdot) = \theta_h^{(2)} Q_\kappa(\mu,\cdot),$$

for every $\mu \in \mathcal{N}(S\times\mathbb{M})$.

Let $\xi_\mu \sim Q_\kappa(\mu,\cdot)$. By Definition 22,

$$\xi_\mu = \sum_{(x,m)\in\mu} \delta_{(x,x)} + \sum_{(x,m_x),(y,m_y)\in\mu}^{x\neq y} E_{\{x,y\}}\,\delta_{(x,y)},$$

where $\{E_{\{x,y\}}\}$ are independent over distinct unordered pairs, with Bernoulli parameter

$$p_{xy} = \kappa(d_S(x,y), m_x, m_y).$$

Applying $\theta_h^{(2)}$,

$$\theta_h^{(2)}\xi_\mu = \sum_{(x,m)\in\mu} \delta_{(hx,hx)} + \sum_{(x,m_x),(y,m_y)\in\mu}^{x\neq y} E_{\{x,y\}}\,\delta_{(hx,hy)}.$$

Since the action is isometric, $d_S(hx,hy) = d_S(x,y)$, hence each shifted edge keeps the same parameter via

$$\kappa(d_S(hx,hy), m_x, m_y) = \kappa(d_S(x,y), m_x, m_y) = p_{xy}.$$

Therefore, $\theta_h^{(2)}\xi_\mu$ has exactly law $Q_\kappa(\theta_h^{(1)}\mu,\cdot)$, proving the claim. Now let $B \in \mathcal{B}(\mathcal{A}(S^2))$. Using

$$C(\gamma,\cdot) = \int Q_\kappa(\mu,\cdot)\, M(\gamma,d\mu),$$

and H-equivariance of M under first-component action,

$$\begin{aligned}
C(\theta_h\gamma, B) &= \int Q_\kappa(\mu,B)\, M(\theta_h\gamma, d\mu) \\
&= \int Q_\kappa(\theta_h^{(1)}\mu, B)\, M(\gamma, d\mu) \\
&= \int \theta_h^{(2)} Q_\kappa(\mu,\cdot)(B)\, M(\gamma,d\mu) \\
&= \int Q_\kappa(\mu, (\theta_h^{(2)})^{-1}B)\, M(\gamma,d\mu) \\
&= C(\gamma, (\theta_h^{(2)})^{-1}B) \\
&= \theta_h^{(2)} C(\gamma,\cdot)(B).
\end{aligned}$$

Hence, $C(\theta_h\gamma,\cdot) = \theta_h C(\gamma,\cdot)$, i.e. C is H-equivariant. ◀

We now provide a proof of Lemma 6.6, which ensures that the limiting object of our sequences is regular (once combined with non-emptiness). The strategy is to control expected edge counts through compact sets via the assumptions, pass to an almost-sure statement on a countable compact exhaustion, and then conclude pointwise fiber-finiteness by choosing a compact set that contains the point.

Proof of Lemma 6.6

Proof. Let $(\Omega, \mathcal{F}, \mathbb{P})$ be a probability space on which all random objects are defined, with

$$\Gamma : \Omega \to \mathcal{N}(S), \qquad \xi : \Omega \to \mathcal{A}(S^2),$$

and $\xi \sim \mathrm{SIRG}(\Gamma, M, \kappa)$. Write the SIRG realization in the form

$$\xi = \sum_{u \in \Gamma} \delta_{(u,u)} + \sum_{\substack{u,v \in \Gamma \\ u \neq v}} E_{\{u,v\}}\, \delta_{(u,v)},$$

where conditional on Γ, each $E_{\{u,v\}}$ is Bernoulli with parameter

$$\kappa(d_S(u,v), W_1, W_2).$$

For $u \neq v$, define

$$p(u,v) := \mathbb{E}[\kappa(d_S(u,v), W_1, W_2)] = \overline{\kappa}(d_S(u,v)\big).$$

Then, by iterated expectation and independence of marks from Γ,

$$\begin{aligned} \mathbb{E}[E_{\{u,v\}} \mid \Gamma] &= \mathbb{E}\big[\mathbb{E}\big[E_{\{u,v\}} \mid \Gamma, W_1, W_2\big] \mid \Gamma\big] \\ &= \mathbb{E}[\kappa(d_S(u,v), W_1, W_2) \mid \Gamma] \\ &= p(u,v) \end{aligned}$$

Fix a compact set $L \subseteq S$, and define the edge count through L:

$$N_L := \sum_{\substack{u,v \in \Gamma \\ u \neq v}} \mathbf{1}_L(u)\, E_{\{u,v\}}.$$

Conditioning on Γ,

$$\mathbb{E}[N_L \mid \Gamma] = \sum_{\substack{u,v \in \Gamma \\ u \neq v}} \mathbf{1}_L(u)\, \overline{\kappa}(d_S(u,v)).$$

Hence, by the defining identity, Equation (27), of the second factorial moment measure $\alpha_\Gamma^{(2)}$,

$$\mathbb{E}[N_L] = \int_{L \times S} \overline{\kappa}(d_S(u,v))\, \alpha_\Gamma^{(2)}(du, dv) < \infty,$$

where the bound comes by assumption. Since $N_L \geq 0$, define for each compact $L \subseteq S$ the event

$$A_L := \{\omega \in \Omega : N_L(\omega) < \infty\}.$$

Then $\mathbb{P}(A_L) = 1$. This is a statement for one compact set at a time. To get a single full-probability event that works simultaneously for all vertices, choose a deterministic countable compact exhaustion $L_n \uparrow S$, which means that

$$\text{each } L_n \text{ is compact, } L_n \subseteq L_{n+1}, \text{ and } \bigcup_{n\in\mathbb{N}} L_n = S.$$

This is possible since S is LC Polish, hence σ-compact. Set

$$\Omega_0 := \bigcap_{n\in\mathbb{N}} A_{L_n} \subseteq \Omega.$$

Then $\mathbb{P}(\Omega_0) = 1$, and for every $\omega \in \Omega_0$ and every n,

$$N_{L_n}(\omega) < \infty.$$

Fix $\omega \in \Omega_0$, and write

$$\gamma := \Gamma(\omega), \qquad \tau := \xi(\omega).$$

Let $x \in S$. If $x \notin \gamma$, then no SIRG atom has endpoint x, so

$$\tau(\{x\} \times S) = \tau(S \times \{x\}) = 0.$$

If $x \in \gamma$, choose n with $x \in L_n$ (possible since $\bigcup_n L_n = S$). Then

$$\sum_{\substack{v\in\gamma \\ v\neq x}} E_{\{x,v\}}(\omega) \leq N_{L_n}(\omega) < \infty.$$

By swapping dummy indices in the definition of N_{L_n}, we also have

$$\sum_{\substack{u,v\in\gamma \\ u\neq v}} \mathbf{1}_{L_n}(u)\, E_{\{u,v\}}(\omega) = \sum_{\substack{u,v\in\gamma \\ u\neq v}} \mathbf{1}_{L_n}(v)\, E_{\{u,v\}}(\omega),$$

so since $x \in L_n$,

$$\sum_{\substack{v\in\gamma \\ v\neq x}} E_{\{x,v\}}(\omega) \leq \sum_{\substack{u,v\in\gamma \\ u\neq v}} \mathbf{1}_{L_n}(v)\, E_{\{u,v\}}(\omega) = N_{L_n}(\omega) < \infty.$$

Also, $(x, x) \in \tau$ contributes only one diagonal atom. Hence,

$$\tau(\{x\} \times S) < \infty, \qquad \tau(S \times \{x\}) < \infty.$$

So every fiber is finite for $\tau = \xi(\omega)$, for all $\omega \in \Omega_0$. Therefore ξ is fiber-finite almost surely, i.e. ξ is a fiber-finite adjacency process. ◀

Lastly, we give a proof of the tightness criterion, Lemma 6.7. The argument combines a crossing-edge tail bound with an induction on graph radius.

Proof of Lemma 6.7

Proof. We prove radial tightness directly from the crossing-edge tail hypothesis in the lemma statement. For $n \in \mathbb{N}$, let Γ_n^x be the Palm-rooted vertex process at x, and let

$$\xi_n^{(x,x)} \sim \mathrm{SIRG}(\Gamma_n^x, M, \kappa_n).$$

Define

$$\bar{\kappa}_n(t) := \mathbb{E}[\kappa_n(t, W_1, W_2)]\,, \qquad O_R := S \setminus B_S(x, R).$$

Fix compact $L \subset S$, $R > 0$, and n, and let $N_n(L, R)$ be the number of oriented SIRG atoms crossing from L to O_R:

$$N_n(L, R) := \sum_{\substack{u,v \in \Gamma_n^x \\ u \neq v}} \mathbf{1}_{\{u \in L,\ v \in O_R\}}\, E_{\{u,v\}}^{(n)},$$

where $E_{\{u,v\}}^{(n)} \in \{0, 1\}$ is the shared Bernoulli edge indicator in the SIRG construction. This construction counts over both orderings, so all crossing edges are covered. Conditioning first on the marked configuration and using that the marking kernel is identically factorable with common mark law ν,

$$\mathbb{E}\Big[E_{\{u,v\}}^{(n)} \mid \Gamma_n^x\Big] = \bar{\kappa}_n(d_{S_n}(u, v))\,.$$

by the same argument as in the prior proof. Hence,

$$\mathbb{E}[N_n(L, R)] = \iint_{L \times O_R} \bar{\kappa}_n(d_{S_n}(u, v))\ \alpha_{\Gamma_n^x}^{(2)}(du, dv).$$

By the tail hypothesis of this lemma, for every compact $L \subset S$ and every $\delta > 0$, there exist $N_*(L, \delta) \in \mathbb{N}$ and $R_*(L, \delta) > 0$ such that for all $n \geq N_*(L, \delta)$,

$$\mathbb{E}[N_n(L, R_*(L, \delta))] < \delta.$$

Therefore, by Markov,

$$\mathbb{P}(N_n(L, R_*(L, \delta)) \geq 1) \leq \delta \quad \text{for all } n \geq N_*(L, \delta).$$

We now show that for every $r \in \mathbb{N}$ and $\varepsilon > 0$, there exist compact $L_r \subset S$ and $N_r \in \mathbb{N}$ such that for all $n \geq N_r$,

$$\mathbb{P}\Big(V_r(\xi_n^{(x,x)}, x) \subset L_r\Big) \geq 1 - \varepsilon,$$

which is exactly radial tightness. We do this using induction. First fix $r = 0$ as the base case. Since the root is fixed at x,

$$V_0(\xi_n^{(x,x)}, x) = \{x\} \quad \text{a.s.},$$

so we may take $L_0 := \{x\}$ and any N_0.

Assume for induction that the claim holds at radius r, and fix $\varepsilon > 0$. By the induction hypothesis with $\varepsilon/2$, there exist compact $L_r \subset S$ and $N_r \in \mathbb{N}$ such that for all $n \geq N_r$,

$$\mathbb{P}\Big(V_r(\xi_n^{(x,x)}, x) \subset L_r\Big) \geq 1 - \frac{\varepsilon}{2}.$$

Applying the crossing estimate above with $L = L_r$ and $\delta = \varepsilon/2$, we obtain $N_* \in \mathbb{N}$ and $R > 0$ such that for all $n \geq N_*$,

$$\mathbb{P}(N_n(L_r, R) \geq 1) \leq \frac{\varepsilon}{2}.$$

At this point we use the compact-containment condition from β-admissibility (condition (3) in Definition 23), which gives that for this R, there exist a compact set $C_R \subset S$ and $N_\beta \in \mathbb{N}$ such that

$$B_{d_{S_n}}(x, R) \subset C_R \quad \text{for all } n \geq N_\beta.$$

Set

$$L_{r+1} := \overline{B_S(x, R)} \cup C_R.$$

By properness of d_S (from Definition 1), $\overline{B_S(x, R)}$ is compact, so L_{r+1} is compact as a union of two compact sets. Now let $n \geq \max\{N_r, N_*, N_\beta\}$. If

$$V_r(\xi_n^{(x,x)}, x) \subset L_r \quad \text{but} \quad V_{r+1}(\xi_n^{(x,x)}, x) \not\subset L_{r+1},$$

then some vertex $y \notin L_{r+1}$ is at graph distance $\leq r+1$ from x. Along a shortest path from x to y, the predecessor of y lies in $V_r(\xi_n^{(x,x)}, x) \subset L_r$ by the inductive hypothesis. Since $y \notin L_{r+1}$, in particular

$$y \notin \overline{B_S(x, R)},$$

hence $y \in O_R$. Therefore, there is at least one crossing edge from L_r to O_R, so

$$\{V_{r+1} \not\subset L_{r+1}\} \subset \{V_r \not\subset L_r\} \cup \{N_n(L_r, R) \geq 1\}.$$

Therefore,

$$\begin{aligned}
\mathbb{P}\Big(V_{r+1}(\xi_n^{(x,x)}, x) \not\subset L_{r+1}\Big) &\leq \mathbb{P}\Big(V_r(\xi_n^{(x,x)}, x) \not\subset L_r\Big) + \mathbb{P}(N_n(L_r, R) \geq 1) \\
&< \frac{\varepsilon}{2} + \frac{\varepsilon}{2} \\
&= \varepsilon.
\end{aligned}$$

So the claim holds at radius $r+1$, and induction is complete. Thus, $\big(\xi_n^{(x,x)}, x\big)_{n \in \mathbb{N}}$ is radially tight. ◀

7 Vertex Processes and Examples

In the last several sections, we have spent time carefully developing a very general framework for determining local weak limits. It is now time to put this machinery to the test, and obtain results for SIRGs under specific vertex processes. This section should not be seen as an exhaustive list of models. Rather, we identify a few clean categories to demonstrate the strength of the framework and the ease with which one can check the conditions of Theorem 6.4 for concrete models. We therefore provide a collection of connection function classes, alongside three chosen vertex process classes (each compatible with all of the connection function classes):

1. Mixed-binomial processes on $\mathbb{R}^d$, converging to Palm Poisson processes.
2. Mixed-binomial processes on the d-dimensional p-adic fields $\mathbb{Q}_p^d$, converging to Palm Poisson processes.
3. Cox processes with stationary directing fields on $\mathbb{R}^d$, converging to Palm Cox with a limiting Palm directing field.

For each, we also give several example use cases. Future extension to other model families, specifically with more complex vertex processes such as Neyman-Scott, determinantal, or Gibbs processes, is very likely possible.

Note that of a grouped LC Polish space (S, H), there are two classifications:

1. If $\lambda_{(S,H)}$ is diffuse, i.e. $\lambda_{(S,H)}(\{x\}) = 0$ for every $x \in S$, we call (S, H) *atomless*.
2. If $\lambda_{(S,H)}$ is not diffuse, we call (S, H) *atomic*.

Recall that $\lambda_{(S,H)}$ is the associated measure on S, as specified in Definition 16. All three of the aforementioned models are natural to atomless spaces, and we enforce this atomlessness in our proposition statements. The spaces $\mathbb{R}^d$ and $\mathbb{Q}_p^d$ are both atomless, and are the canonical choices for atomless LC Polish spaces that admit β-admissable approximations.

This section proceeds as follows. First, Section 7.1 records the shared context and standing assumptions used throughout the model constructions. Next, in Section 7.2, we introduce the two classes of point processes used in our examples, namely mixed binomial and Cox point processes. Then, Section 7.3 develops our first convergence paradigm, in which mixed binomial vertex processes converge to Poisson limits on $\mathbb{R}^d$ and $\mathbb{Q}_p^d$, together with the corresponding connection kernel classes. Finally, Section 7.4 treats our second paradigm, where Cox vertex processes converge to Cox limits under analogous kernel assumptions.

7.1 Context and Assumptions Overview

Before moving forward, we isolate a shared context for the model constructions in this section. It is natural to do this here, rather than earlier in the thesis, because at this stage the ambient setup has become quite standard across the examples. The same basic ingredients recur, and the main variation lies in how

we specialize them. Earlier results typically differed slightly in their underlying hypotheses or objects, so introducing a single standing context there would have obscured those distinctions rather than clarified them. Here, by contrast, a common context makes the bookkeeping cleaner and avoids repeatedly restating the same structural assumptions in each proposition, lemma, and corollary.

Conceptually, one should think of the context below as specifying a family of admissible tuples. To *specialize* the context in a later statement means to impose additional conditions by fixing some of the entries of that tuple, or by requiring them to belong to a more specific class. In effect, the specialization is simply appended to the context as an extra constraint. It is possible in principle that a given specialization leaves no admissible tuple at all, so that the specialized context is empty. In most of our examples, however, we do not prove non-emptiness as a separate preliminary step, because it is usually built directly into the construction. See the context below.

Context 1 SIRG Context

Consider a tuple

$$((S,H)^2, (S_n,H_n)^2_{n\in\mathbb{N}}, x, \nu, (\kappa_n)_{n\in\mathbb{N}}, \kappa_\infty, (\Gamma_n)_{n\in\mathbb{N}}, \Gamma_\infty, (\xi_n)_{n\in\mathbb{N}}, \xi_\infty).$$

This tuple is called an *SIRG Context* when it satisfies the following conditions:

1. $(S,H)^2$ is a square grouped LC Polish space;
2. $x \in S$ and $(S_n,H_n)^2 \subset_{\beta,x} S^2$;
3. ν is a probability measure on some LC Polish space $\mathbb{M}$;
4. $(\kappa_n)_{n\in\mathbb{N}}, \kappa_\infty$ are connection functions mapping $\mathbb{R}^+ \times \mathbb{M} \times \mathbb{M} \to [0,1]$;
5. For each $n \in \mathbb{N}$, Γ_n is a point process on (S_n,H_n), and Γ_∞ is a point process on S;
6. For each $n \in \mathbb{N}$, $\xi_n \sim \mathrm{SIRG}(\Gamma_n, M, \kappa_n)$ is a SIRG[a] on $(S_n,H_n)^2$, taking marks in $\mathbb{M}$, where M is an identically factorable marking kernel with shared law ν;
7. $\xi_\infty \sim \mathrm{SIRG}(\Gamma_\infty, M, \kappa_\infty)$ is a SIRG on S^2 taking marks in $\mathbb{M}$.

[a] Naturally these are equivariant since the marking kernel is identically factorable and β-admissibility gives isometric actions.

The assumptions of Theorem 6.4 can then be expressed as follows, and we will henceforth in this section refer to these in most cases rather than the original theorem to make the upcoming chaining of assumptions clean. We also will shortly provide a clear corollary mimicking Theorem 6.4 with these assumptions.

▶ **Assumption 7.1.** *Given an instance of Context 1, assume that Γ_n is stationary on (S_n,H_n) for each $n \in \mathbb{N}$, and that $\Gamma_n^x \Rightarrow \Gamma_\infty$.*

▶ **Assumption 7.2.** *Given an instance of Context 1, assume that $\kappa_n \to^s \kappa_\infty$.*

▶ **Assumption 7.3.** *Given an instance of Context 1, assume that κ_∞ is $D(\Gamma_\infty)\otimes \nu\otimes\nu$-a.s. continuous.*

▶ **Assumption 7.4.** *Given an instance of Context 1, assume that condition (3) of Theorem 6.4 holds.*

▶ **Assumption 7.5.** *Given an instance of Context 1, assume that conditions (4-5) of Theorem 6.4 hold.*

With this notation in place, Theorem 6.4 can be repackaged as a short consequence of the context and the five assumptions above. This is precisely the convenience of introducing the shared context here. Once the ambient tuple has been fixed, later arguments can focus on verifying the few properties that actually vary from model to model.

▶ **Corollary 7.6.** *Let the tuple*

$$((S,H)^2,(S_n,H_n)^2_{n\in\mathbb{N}},x,\nu,(\kappa_n)_{n\in\mathbb{N}},\kappa_\infty,(\Gamma_n)_{n\in\mathbb{N}},\Gamma_\infty,(\xi_n)_{n\in\mathbb{N}},\xi_\infty)$$

be an instance of Context 1 satisfying Assumptions 7.1–7.5. Then,

$$G(\xi_n,U_n)\Rightarrow G(\xi_\infty,x),\qquad U_n\sim\mathbb{U}(\xi_n,\cdot).$$

Proof. Naturally, Assumption 7.1 makes all $(\xi_n)_{n\in\mathbb{N}}$ stationary and equivariant, by Lemma 6.3 applied with the identically factorable marking kernel M with shared law ν, which is intrinsically equivariant. Thus, we match the context of Theorem 6.4. Then, Assumptions 7.1–7.5 match the conditions of Theorem 6.4 verbatim. ◀

Since we care about atomless spaces, we typically impose one further assumption guaranteeing that the chosen subspaces genuinely exhaust the ambient space. This is mainly useful for the convergence results we want to derive, but it is also the standard way these examples are naturally constructed. We phrase this additional requirement using the reference measure on the ambient space, which we recall is attached by Definition 16.

▶ **Assumption 7.7.** *Given an instance of Context 1, assume that*

$$\lambda_{(S,H)}|_{S_n}\to\lambda_{(S,H)}$$

vaguely (against continuous compactly supported test functions).

7.2 Mixed Binomial and Cox Point Processes

In this section, we will define the two classes of point processes that we use among all of our examples. These are the classes of *mixed binomial point processes* and *Cox point processes*. The former are the simpler of the two, and are in fact one of the simplest examples in point process theory, and thus we begin there.

A mixed binomial vertex process $\Gamma \sim \mathrm{MB}(N, \mu)$ on an atomless grouped LC Polish space (S, H) is a point process given by

$$\Gamma :\overset{d}{=} \sum_{i=1}^{N} \delta_{X_i}, \tag{30}$$

where N is an almost surely finite random natural number (including zero) and the family $(X_i)_{i\in\mathbb{N}}$ are i.i.d. S-valued random variables with shared law μ, all independent of N. In order to keep Γ simple, it must hold that μ is diffuse, which is the convention that we carry henceforth. Effectively, this model represents placing a random number of vertices in space independently in the same random way. They are, however, limited in their representational ability, as they must remain finite, and thus are most effective on compact spaces. It is for this reason that they are perfect as our baseline models on a sequence of compact β-admissible subspaces. We will touch more on their role in sequences in the next section, but for now it is useful to describe their behavior under the Palm transformation, as this gives some intuition behind their properties. For Γ, as given in Equation (30), it holds for Λ_Γ-a.e. $x \in S$ that[13]

$$\Gamma^x = \delta_x + \sum_{i=1}^{N^*-1} \delta_{X_i} \tag{31}$$

whenever $0 < \mathbb{E}[N] < \infty$, where N^* is the size-biased version of N, satisfying the functional identity

$$\mathbb{E}[f(N^*)] = \frac{\mathbb{E}[Nf(N)]}{\mathbb{E}[N]} \tag{32}$$

for every bounded measurable $f : \mathbb{R}^+ \to \mathbb{R}$. These exist by [17, Eq. (1.4.18)]. The simplicity of this Palm version highlights the strong regularity of these models. Specifically, the Palm version of a mixed binomial point process is useful because it gives an especially transparent interpretation of what it means to observe the process from the perspective of one of its own points. Conditioning on seeing a point at x does not distort the spatial distribution of the other vertices. It simply places one vertex at x, and then samples the remaining vertices independently of the original law μ. The only change is that the total number of vertices is size biased.

This is exactly what one should expect from a uniformly chosen vertex. Larger configurations are more likely to be seen, because they contain more possible choices for the root. In this sense, mixed binomial processes provide a clean baseline, as rooting changes the count distribution, but leaves the independent spatial placement mechanism intact.

It is also useful to give the closed forms of intensities and second factorial moments of these Palm variants, as these are specifically useful for satisfying conditions (4-5) of Theorem 6.4. More specifically, we give these for the reduced

[13] This is presented as a standard fact. One can use the Palm identity from Sec. 5.1 to verify this claim.

Palm variants. For a point process η, its reduced Palm variant at x, denoted $\eta^{!x}$, is simply $\eta^x - \delta_x$, under the same existence conditions. We use these variants because they have a cleaner closed form, and with some minor variation to the aforementioned conditions can still be used. For this we have the following lemma.

▶ **Lemma 7.8.** *Consider a grouped LC Polish space (S, H) and let $\Gamma \sim \mathrm{MB}(N, \mu)$ where N is an almost surely finite random natural number and the family $(X_i)_{i\in\mathbb{N}}$ are i.i.d. S-valued random variables with shared diffuse law μ, all independent of N. Then, for any $x \in S$ for which Γ^x exists, and Borel $A, B \subset S$, it holds that*

$$\Lambda_{\Gamma^{!x}}(A) = \frac{\mathbb{E}[N(N-1)]}{\mathbb{E}[N]}\mu(A),$$

$$\alpha^{(2)}_{\Gamma^{!x}}(A \times B) = \frac{\mathbb{E}[N(N-1)(N-2)]}{\mathbb{E}[N]}\mu(A)\mu(B).$$

Proof. Let $A, B \subseteq S$ be Borel. Then,

$$\begin{aligned}
\Lambda_{\Gamma^{!x}}(A) &= \mathbb{E}[\Gamma^{!x}(A)] && \text{Eq. (9)}\\
&= \mathbb{E}\left[\sum_{i=1}^{N^*-1} \mathbf{1}_A(X_i)\right] && \text{Eq. (31)}\\
&= \mathbb{E}[N^* - 1]\mu(A) && \text{Eq. (30)}\\
&= \frac{\mathbb{E}[N(N-1)]}{\mathbb{E}[N]}\mu(A), && \text{Eq. (32)}
\end{aligned}$$

giving the intensity measure claim. Similarly for the second factorial moment measure,

$$\begin{aligned}
\alpha^{(2)}_{\Gamma^{!x}}(A \times B) &= \mathbb{E}\left[\sum_{u,v\in\Gamma^{!x}}^{u\neq v} \mathbf{1}_A(u)\mathbf{1}_B(v)\right] && \text{Eq. (26)}\\
&= \mathbb{E}\left[\sum_{1\leq i,j\leq N^*-1}^{i\neq j} \mathbf{1}_A(X_i)\mathbf{1}_B(X_j)\right] && \text{Eq. (31)}\\
&= \mathbb{E}[(N^* - 1)(N^* - 2)]\mu(A)\mu(B) && \text{Eq. (30)}\\
&= \frac{\mathbb{E}[N(N-1)(N-2)]}{\mathbb{E}[N]}\mu(A)\mu(B), && \text{Eq. (32)}
\end{aligned}$$

which concludes the proof. ◀

This serves as a sufficient introduction to mixed binomial point processes, and thus we shift to the notion of a Cox process. For illumination, we start by presenting a simpler variant, the *Poisson point process*. However, it may first be useful to reflect on diffuse versus atomic measures.

In the mixed binomial example, we saw that the distribution with which we placed each independent point needed to be diffuse. At a high level, this makes sense, as if $\mu(z) > 0$ for some $z \in S$, it would immediately follow that $\mathbb{P}(\Gamma(z) > 1)$ has a non-zero probability of occurring, which would cause the process to not be simple. What is important to note, however, is that this is simply a single instance of a larger phenomenon, which is that in order to ensure that a point process η is simple, it is often very beneficial, and sometimes necessary, to require Λ_η to be diffuse.

Furthermore, while we have thus far treated the intensity measure as something derived from a point process, this perspective is often flipped on its head, where one instead takes a diffuse intensity measure Φ first, and constructs a point process satisfying that intensity. We use Φ here instead of Λ to disambiguate this change in perspective. The most natural construction of this is the aforementioned Poisson point process. Given a boundedly finite[14] diffuse Borel[15] measure Φ, a Poisson point process $\Gamma \sim \mathrm{PPP}(\Phi)$ is the unique point process satisfying

$$\Gamma(A) \sim \mathrm{Poi}(\Phi(A))$$

for every bounded Borel set $A \subseteq S$, and such that $\Gamma(A_1), \ldots, \Gamma(A_k)$ are independent whenever $A_1, \ldots, A_k \subseteq S$ are pairwise disjoint bounded Borel sets. Here, Poi is the classical Poisson count given for $\lambda \geq 0$ by

$$\mathrm{Poi}(\lambda)(\{k\}) = e^{-\lambda}\frac{\lambda^k}{k!}, \qquad k \in \mathbb{N}.$$

By construction, $\Lambda_\Gamma = \Phi$. Moreover, because Φ is diffuse and boundedly finite, this process is simple and boundedly finite almost surely.

Following from this, the Cox process generalizes the Poisson process to the case where Φ is random, in which case it is called a *directing measure*. In order to properly state this, we define the space of boundedly finite Borel measures on S, denoted $\mathcal{M}(S)$, analogously to the way we constructed $\mathcal{N}(S)$, also carrying the vague topology. Indeed, $\mathcal{N}(S) \subset \mathcal{M}(S)$, carrying the integer-valued condition. In this space, we can measurably define directing measures to be random elements in $\mathcal{M}(S)$ in the same way we constructed point process as random elements in $\mathcal{N}(S)$.

Hence, given a diffuse[16] random directing measure Φ, the Cox process $\Gamma \sim \mathrm{Cox}(\Phi)$ is the unique point process for which, conditionally on Φ, it holds that $\Gamma \sim \mathrm{PPP}(\Phi)$. Similarly, the resulting point process is also boundedly finite and simple. In the Euclidean case used below, stationarity of the directing measure implies stationarity of the Cox process by [32].

As for the Palm distribution, it turns out that the Palm variant of a Cox process is closely related to the Cox process on the Palm variant of its directing measure.

[14] This is an identical to Def. 2 for diffuse measures.

[15] Measure defined on the Borel σ-algebra, standard in this context.

[16] Here, this means almost surely diffuse. Boundedly finite is not restated, left to the context of $\mathcal{M}(S)$.

Specifically, we can take the Palm versions of directing measures analogously to in Definition 18, via the identity

$$\mathbb{E}\left[\int_S f(y,\Phi)\,\Phi(dy)\right] = \int_S \mathbb{E}[f(x,\Phi^x)]\,\Lambda_\Phi(dx), \tag{33}$$

where the original sum on the left-hand side is replaced by an integral. It also becomes especially clear here why we chose to use the notation Φ instead of Λ, as it allows us to cleanly denote intensities of directing measures as well. Then, by [37, Thm. 8], it is established that for $\Gamma \sim \mathrm{Cox}(\Phi)$, $\eta \sim \mathrm{Cox}(\Phi^x)$, it holds that

$$\Gamma^x \overset{d}{=} \delta_x + \eta \tag{34}$$

under the standard existence conditions. For intensities and moments of these Palm variants, we give a lemma analogous to Lemma 7.8.

▶ **Lemma 7.9.** *Consider a grouped LC Polish space (S,H), and let $\Gamma \sim \mathrm{Cox}(\Phi)$, where Φ is a diffuse directing measure on S. Then, for any $x \in S$ for which the Palm version Φ^x exists, and Borel $A, B \subset S$, it holds that*

$$\Lambda_{\Gamma^{!x}}(A) = \mathbb{E}[\Phi^x(A)],$$

$$\alpha^{(2)}_{\Gamma^{!x}}(A \times B) = \mathbb{E}[\Phi^x(A)\Phi^x(B)].$$

Proof. From Equation (34) we obtain

$$\Gamma^{!x} \mid \Phi^x \sim \mathrm{PPP}(\Phi^x).$$

Conditional on Φ^x, the reduced Palm process is therefore an inhomogeneous Poisson point process with directing measure Φ^x. Recall that the defining first and second factorial moment measures of a Poisson point process are

$$\Lambda_{\Gamma^{!x}|\Phi^x}(A) = \Phi^x(A),$$

$$\alpha^{(2)}_{\Gamma^{!x}|\Phi^x}(A \times B) = \Phi^x(A)\Phi^x(B).$$

Taking expectations and applying iterated expectation gives

$$\Lambda_{\Gamma^{!x}}(A) = \mathbb{E}[\Lambda_{\Gamma^{!x}|\Phi^x}(A)] = \mathbb{E}[\Phi^x(A)],$$

$$\alpha^{(2)}_{\Gamma^{!x}}(A \times B) = \mathbb{E}[\alpha^{(2)}_{\Gamma^{!x}|\Phi^x}(A \times B)] = \mathbb{E}[\Phi^x(A)\Phi^x(B)],$$

as claimed. ◀

With the basics of these processes established, we move on to our first paradigm, where we use mixed binomial vertex processes on a β-admissible sequence of compact spaces to arrive at a Poisson limiting vertex process in the ambient space.

7.3 Mixed Binomial Vertex Processes with Poisson Limits

7.3.1 General Theory for Poisson Limits

First fix for each $n \in \mathbb{N}$,

$$\Gamma_n \sim \mathrm{MB}(N_n, \mu_n),$$

where $(N_n)_{n\in\mathbb{N}}$ is a sequence of random counts and for each $n \in \mathbb{N}$, μ_n is a sequence of probability measures on S_n. Assume further that $N_n \to \infty$ in probability, as otherwise our limiting object is finite and does not require our theory.

In general, the weak limit of the sequence $(\Gamma_n)_{n\in\mathbb{N}}$ is $\Gamma_\infty \sim \mathrm{Cox}(\Phi)$ whenever

$$N_n\mu_n \Rightarrow \Phi \tag{35}$$

in the vague topology on $\mathcal{M}(S)$. This is given by [19, Thm. 30.5]. However, since we want to apply Theorem 6.4 to this sequence and its limit, the elements of the sequence and the possible limiting distributions are heavily restricted by our required conditions, and consequently we can not exploit this result in its full generality.

Firstly, we need each Γ_n to be stationary on (S_n, H_n). Due to the factorable nature of mixed binomial point processes, this is only possible when $X \sim \mu_n$ is uniform on S_n for each $n \in \mathbb{N}$. The standard way to achieve this is to define for each $n \in \mathbb{N}$ the probability measure

$$\mu_n = \frac{1}{\lambda_{(S,H)}(S_n)}\lambda_{(S,H)}|_{S_n}, \tag{36}$$

where we recall that $\lambda_{(S,H)}$ is the ambient reference measure on S fixed by Definition 16. This alone restricts the class of possible limiting directing measures to measures of the form

$$\Phi = W\lambda_{(S,H)},$$

where W is a non-negative real random variable. Under Assumption 7.7, this is satisfied whenever

$$\frac{N_n}{\lambda_{(S,H)}(S_n)} \Rightarrow W, \tag{37}$$

by Equations (35) and (36). However, we also must take into account the vertex concentration condition, condition (3), of Theorem 6.4, which in this context can be expressed as

$$\mathbb{E}\left[\left|\frac{N_n}{\mathbb{E}[N_n]} - 1\right|\right] \to 0.$$

This further forces W to be a deterministic constant $c > 0$, since

$$\frac{N_n}{\lambda_{(S,H)}(S_n)} = \frac{N_n}{\mathbb{E}[N_n]} \cdot \frac{\mathbb{E}[N_n]}{\lambda_{(S,H)}(S_n)} \;\Rightarrow\; W = c, \tag{38}$$

and the first factor goes to 1 via concentration while the second is a deterministic sequence. So, all possible limits in this context are given by $\Gamma \sim \mathrm{PPP}(c\lambda_{(S,H)})$. It also turns out that under a third moment condition on this ratio, we can transfer this convergence to the Palm variants. See the following assumption and associated lemma, which handle stationarity convergence of the Palm processes via stability of the counts.

▶ **Assumption 7.10.** *Given an instance of Context 1, some $c \in \mathbb{R}^+$, and some count sequence $(N_n)_{n\in\mathbb{N}}$, assume that $N_n \to \infty$ in probability, and that for some constant $N_0 \in \mathbb{N}$,*

$$0 < \mathbb{E}[N_n] < \infty \text{ for all } n \geq N_0,$$

$$\frac{N_n}{\lambda_{(S,H)}(S_n)} \Rightarrow c, \qquad \sup_{n\geq N_0} \mathbb{E}\left[\left(\frac{N_n}{\lambda_{(S,H)}(S_n)}\right)^3\right] < \infty.$$

▶ **Lemma 7.11.** *Let the tuple*

$$((S,H)^2, (S_n,H_n)^2_{n\in\mathbb{N}}, x, \nu, (\kappa_n)_{n\in\mathbb{N}}, \kappa_\infty, (\Gamma_n)_{n\in\mathbb{N}}, \Gamma_\infty, (\xi_n)_{n\in\mathbb{N}}, \xi_\infty)$$

be an instance of Context 1 specializing $(\Gamma_n)_{n\in\mathbb{N}}, \Gamma_\infty$ to

$$\Gamma_n \sim \mathrm{MB}(N_n, \mu_n) \qquad \text{for all } n \in \mathbb{N},$$

$$\Gamma_\infty \overset{d}{=} \delta_x + \eta, \qquad \eta \sim \mathrm{PPP}(c\lambda_{(S,H)}),$$

where $(N_n)_{n\in\mathbb{N}}$ is a sequence of random natural numbers, $c \in \mathbb{R}^+$, and $(\mu_n)_{n\in\mathbb{N}}$ is a sequence of measures given by

$$\mu_n = \frac{1}{\lambda_{(S,H)}(S_n)} \lambda_{(S,H)}|_{S_n}.$$

Then, Assumption 7.7 and Assumption 7.10 together give Assumption 7.1.

Proof. Recall throughout that transitivity is given by Context 1. We first show stationarity of the elements of the sequence. Fix $n \in \mathbb{N}$. Then, for each $h \in H_n$, stationarity of uniform random variables gives

$$\Gamma_n \overset{d}{=} \sum_{i=1}^{N_n} \delta_{X_i} \overset{d}{=} \sum_{i=1}^{N_n} \delta_{X_i} \overset{d}{=} \theta_h \Gamma_n.$$

Since n was arbitrary we conclude stationarity throughout the sequence. Now the convergence. By Equation (31) and the fact that $0 < \mathbb{E}[N_n] < \infty$ eventually via Assumption 7.10, all Palm versions $(\Gamma_n^x)_{n\geq N_1}$ exist uniquely for some $N_1 \in \mathbb{N}$. Since deleting finitely many initial terms does not affect weak convergence, it suffices to show that for $\eta_n \sim \mathrm{MB}(N_n^* - 1, \mu_n)$, it holds that $\eta_n \Rightarrow \eta$ for all sufficiently large n. By Equation (37) and Assumption 7.7, this follows from

$$\frac{N_n^* - 1}{\lambda_{(S,H)}(S_n)} \Rightarrow c,$$

which is what we now show. Write for notational shorthand

$$Y_n := \frac{N_n}{\lambda_{(S,H)}(S_n)}.$$

Then, Assumption 7.10 gives $Y_n \Rightarrow c$, and implies that $(Y_n)_{n>N_0}$ is uniformly integrable for some $N_0 \in \mathbb{N}$. This means that

$$\lim_{K\to\infty} \sup_{n>N_0} \mathbb{E}[Y_n \mathbf{1}_{\{Y_n>K\}}] = 0.$$

From uniform integrability, it holds that for every continuous function $h : \mathbb{R}^+ \to \mathbb{R}$ that grows at most linearly,

$$\mathbb{E}[h(Y_n)] \to \mathbb{E}[h(c)],$$

since $Z_n = h(Y_n)$ is uniformly integrable as well. Hence, for every bounded continuous $g : \mathbb{R}^+ \to \mathbb{R}$, $h(y) := yg(y)$ grows at most linearly, and Equation (32) applied for all sufficiently large n with $f_n(k) := g(k/\lambda_{(S,H)}(S_n))$ yields

$$\mathbb{E}\left[g\left(\frac{N_n^*}{\lambda_{(S,H)}(S_n)}\right)\right] = \frac{\mathbb{E}[Y_n g(Y_n)]}{\mathbb{E}[Y_n]} = \frac{\mathbb{E}[h(Y_n)]}{\mathbb{E}[Y_n]} \to \frac{cg(c)}{c} = g(c),$$

and therefore Proposition 2.1 gives

$$\frac{N_n^*}{\lambda_{(S,H)}(S_n)} \Rightarrow c.$$

Since $N_n \to \infty$ in probability by Assumption 7.10, this also forces divergence of $\lambda_{(S,H)}(S_n)$. Therefore,

$$\frac{N_n^* - 1}{\lambda_{(S,H)}(S_n)} = \frac{N_n^*}{\lambda_{(S,H)}(S_n)} - \frac{1}{\lambda_{(S,H)}(S_n)} \Rightarrow c.$$

◀

Thus, Assumption 7.1 is covered. Intuitively, it ensures that the spatial volume scales proportionally to the number of atoms, while the additional moment conditions guarantee that the Palm versions exist and that the vertex counts do not vary wildly. In fact, it turns out that this is sufficient to conclude Assumption 7.4 as well. See the following.

▶ **Lemma 7.12.** *Let the tuple*

$$((S,H)^2, (S_n, H_n)^2_{n\in\mathbb{N}}, x, \nu, (\kappa_n)_{n\in\mathbb{N}}, \kappa_\infty, (\Gamma_n)_{n\in\mathbb{N}}, \Gamma_\infty, (\xi_n)_{n\in\mathbb{N}}, \xi_\infty)$$

be an instance of Context 1 specializing $(\Gamma_n)_{n\in\mathbb{N}}, \Gamma_\infty$ *to*

$$\Gamma_n \sim \mathrm{MB}(N_n, \mu_n) \qquad \textit{for all } n \in \mathbb{N},$$

$$\Gamma_\infty \overset{d}{=} \delta_x + \eta, \qquad \eta \sim \mathrm{PPP}(c\lambda_{(S,H)}),$$

where $(N_n)_{n\in\mathbb{N}}$ *is a sequence of random natural numbers,* $c \in \mathbb{R}^+$*, and* $(\mu_n)_{n\in\mathbb{N}}$ *is a sequence of measures given by*

$$\mu_n = \frac{1}{\lambda_{(S,H)}(S_n)} \lambda_{(S,H)}|_{S_n}.$$

Then, Assumption 7.7 and Assumption 7.10 together give Assumption 7.4.

Proof. Write again for notational shorthand

$$Y_n := \frac{N_n}{\lambda_{(S,H)}(S_n)}.$$

By Assumption 7.10 and Assumption 7.7, $Y_n \Rightarrow c$, and further gives that both $(Y_n)_{n>N_0}$ and $(Y_n^2)_{n>N_0}$ are uniformly integrable. Hence,

$$\mathbb{E}[Y_n] \to c, \qquad \mathbb{E}[Y_n^2] \to c^2,$$

and therefore $\mathrm{Var}(Y_n) \to 0$. Since $\mathbb{E}[Y_n] \to c > 0$,

$$\mathbb{E}\left[\left|\frac{N_n}{\mathbb{E}[N_n]} - 1\right|\right] = \mathbb{E}\left[\left|\frac{Y_n}{\mathbb{E}[Y_n]} - 1\right|\right] = \frac{\mathbb{E}[|Y_n - \mathbb{E}[Y_n]|]}{\mathbb{E}[Y_n]} \leq \frac{\sqrt{\mathrm{Var}(Y_n)}}{\mathbb{E}[Y_n]} \to 0.$$

where the bound comes from the Cauchy-Schwarz inequality. This is exactly condition (3). ◀

It remains only to verify fiber-finiteness and radial tightness, via Assumption 7.5. In this paradigm, the key determining factor is the decay rate of the connection functions. It suffices to show that the connection kernels eventually uniformly decay at a rate that is inversely proportional to the rate at which the volume of a closed ball of radius r grows in r. This is measured by the radial growth function

$$V_x(t) := \lambda_{(S,H)}(B_S(x,t)).$$

In the spaces we will work with, this happens to be a polynomial decay condition related to the dimension of the space. However, we state the assumption in the general case, as it better illuminates the true determining condition. See the assumption below.

▶ **Assumption 7.13.** *Given a sequence of connection functions* $(\kappa_n)_{n\in\mathbb{N}}$*, assume that there exist* $N_0 \in \mathbb{N}$*,* $t_0 > 0$*, and a measurable non-increasing function* $\phi : \mathbb{R}^+ \to \mathbb{R}^+$ *such that*

$$\sup_{n\geq N_0} \overline{\kappa}_n(t) \leq \phi(t) \quad \textit{for all } t \geq t_0,$$

where ϕ *satisfies*

$$\int_0^\infty \phi(t)\, dV_x(t) < \infty.$$

The integral here is a Stieltjes integral. The notation denotes integrating ϕ against the measure μ_x induced by the non-decreasing function V_x via $\mu_x((a,b]) := V_x(b) - V_x(a)$, which records radial-volume increments around x.

As such, the increment $dV_x(t)$ represents how much reference-measure mass is added when radius increases at distance t, so $\phi(t)\, dV_x(t)$ is the corresponding upper bound on distance-t connection contributions. Since this is integrable over $[0,\infty)$, the total contribution from sufficiently large distances can be made arbitrarily small, uniformly for large n. This is exactly the control needed for condition (5), and the same bound also yields finiteness of the full $L \times S$ integral in condition (4) for the limit process. See the following lemma.

▶ **Lemma 7.14.** *Let the tuple*

$$((S,H)^2, (S_n,H_n)^2_{n\in\mathbb{N}}, x, \nu, (\kappa_n)_{n\in\mathbb{N}}, \kappa_\infty, (\Gamma_n)_{n\in\mathbb{N}}, \Gamma_\infty, (\xi_n)_{n\in\mathbb{N}}, \xi_\infty)$$

be an instance of Context 1 specializing $(\Gamma_n)_{n\in\mathbb{N}}, \Gamma_\infty$ *to*

$$\Gamma_n \sim \mathrm{MB}(N_n, \mu_n) \qquad \text{for all } n \in \mathbb{N},$$

$$\Gamma_\infty \overset{d}{=} \delta_x + \eta, \qquad \eta \sim \mathrm{PPP}(c\lambda_{(S,H)}),$$

where $(N_n)_{n\in\mathbb{N}}$ *is a sequence of random natural numbers,* $c \in \mathbb{R}^+$*, and* $(\mu_n)_{n\in\mathbb{N}}$ *is a sequence of measures given by*

$$\mu_n = \frac{1}{\lambda_{(S,H)}(S_n)} \lambda_{(S,H)}|_{S_n}.$$

Then, Assumptions 7.7, 7.10 and 7.13 jointly give Assumption 7.5.

The proof is somewhat technical and not as illuminating as the prior two, and thus it is deferred to Appendix A.6.

We now have enough to apply Theorem 6.4. Thus, we give the general proposition governing this convergence paradigm below, and a proof using the previous lemmas.

▶ **Proposition 7.15** Local Weak Convergence to Palm Poisson Limits. *Let the tuple*

$$((S,H)^2, (S_n,H_n)^2_{n\in\mathbb{N}}, x, \nu, (\kappa_n)_{n\in\mathbb{N}}, \kappa_\infty, (\Gamma_n)_{n\in\mathbb{N}}, \Gamma_\infty, (\xi_n)_{n\in\mathbb{N}}, \xi_\infty)$$

be an instance of Context 1 specializing $(\Gamma_n)_{n\in\mathbb{N}}, \Gamma_\infty$ *to*

$$\Gamma_n \sim \mathrm{MB}(N_n, \mu_n) \qquad \text{for all } n \in \mathbb{N},$$

$$\Gamma_\infty \overset{d}{=} \delta_x + \eta, \qquad \eta \sim \mathrm{PPP}(c\lambda_{(S,H)}),$$

where $(N_n)_{n\in\mathbb{N}}$ *is a sequence of random natural numbers,* $c \in \mathbb{R}^+$*, and* $(\mu_n)_{n\in\mathbb{N}}$ *is a sequence of measures given by*

$$\mu_n = \frac{1}{\lambda_{(S,H)}(S_n)} \lambda_{(S,H)}|_{S_n}.$$

Then, under Assumptions 7.2, 7.3, 7.7, 7.10 and 7.13, it holds that

$$G(\xi_n, U_n) \Rightarrow G(\xi_\infty, x)$$

where $U_n \sim \mathbb{U}(\xi_n, \cdot)$.

Proof. We use Corollary 7.6. Thus, we need to verify Assumptions 7.1–7.5. We do this one by one below:

Assumption 7.1: This is given by Lemma 7.11 under Assumptions 7.7 and 7.10.

Assumption 7.2 This remains assumed directly.

Assumption 7.3 This remains assumed directly.

Assumption 7.4: This is given by Lemma 7.12 under Assumptions 7.7 and 7.10.

Assumption 7.5: This is given by Lemma 7.14 under Assumptions 7.7, 7.10 and 7.13.

This concludes the proof via application of Corollary 7.6. ◀

At this stage, these types of arguments are largely verification-based. As the general convergence framework and supporting lemmas are in place, the proofs consist primarily of checking that the stated assumptions satisfy the required conditions. Consequently, these proofs are necessarily somewhat bland, but their role is mainly organizational rather than conceptual.

This set of assumptions is finally at a level where one can easily instantiate and check a real world model. With the theory for this paradigm complete, we will now take some time to give material examples and place it in the literature, to demonstrate how these checks typically go.

7.3.2 Poisson Limit Examples in $\mathbb{R}^d$

The prototypical example of a model to which one would want to apply this result is the binomial process blow-up model, as in [18]. There, the finite spaces are Euclidean boxes or, equivalently for local purposes, n-volume flat tori $\mathbb{T}^d_n \subset_{\beta,0} \mathbb{R}^d$, equipped with the torus metric

$$d_{\mathbb{T}^d_n}(x, y) = \sqrt{\left(\sum_{i=1}^{d} \min\left\{ |x_i - y_i|, n^{1/d} - |x_i - y_i| \right\}^2 \right)}$$

acting on themselves via translation. These are tori with total volume n, not side length n, i.e.

$$\mathbb{T}^d_n := \left[-\frac{n^{1/d}}{2}, \frac{n^{1/d}}{2} \right)^d.$$

The sequence of vertex processes then places n points in each corresponding torus, and the limiting Palm vertex process is

$$\Gamma_\infty = \delta_0 + \eta, \qquad \eta \sim \mathrm{PPP}(\lambda_{\mathbb{R}^d}),$$

where $\lambda_{\mathbb{R}^d}$ is the standard Lebesgue measure. We fix this measure to be canonical for $\mathbb{R}^d$ as a grouped LC Polish space acting on itself, and fix its restriction to $\mathbb{T}_n^d$ to be canonical as well, when treating $\mathbb{T}_n^d$ as a grouped LC Polish space acting on itself. The local limit theorem of [18, Thm. 1.9] then yields convergence to the corresponding infinite SIRG. Our Proposition 7.15 recovers this vertex-process limit at a comparable level of generality, but with some differences:

1. [18] use these subspaces with the real metric, not the torus metric. However, in the limit, the differences between these metrics balance out, and our models give the same local limits;
2. our result gives convergence in distribution of the induced rooted graph, whereas [18] also proves convergence locally in probability;
3. their mark assumption is stated at the level of empirical convergence of weights to an identically factorable marking kernel, while our present formulation assumes an identically factorable marking kernel for all n;
4. the present theorem states the required connection function convergence in a sequential form. Examples from [18] whose verification is only given pointwise therefore require an additional sequential-convergence check before they can be imported directly into the present framework. However, our model can cover all of the same types of limiting connection functions.

Thus, our model can also handle PSIRGs, GIRGs, threshold HRGs, parametrized HRGs, continuum scale-free percolation models, weight-dependent random connection models, etc. under analogous constructions to in [18] and the aforementioned differing assumptions. For more detail on these models, see Appendix B.2, although as these results are not novel, we do not give a considerable exposition on their intuition or use cases. For this it is advised that the reader see [18] and the referenced material there.

It is worth noting that in this $\mathbb{R}^d$ setting, the analogous decay assumption in [18, Ass. 1.8] is expressed as

$$\overline{\kappa}_\infty(t) \leq t^{-\alpha}$$

eventually in t, where $\alpha > d$. This appears considerably less complex than our Assumption 7.13, but in fact we can reduce to a similar form. Specifically, it holds that in $\mathbb{R}^d$,

$$V_x(t) \asymp t^d,$$

i.e. $c_1 t^d \leq V_x(t) \leq c_2 t^d$ for some $c_1, c_2 \in \mathbb{R}^+$. Thus, Assumption 7.13 resolves to the following assumption.

▶ **Assumption 7.16.** *Given a sequence of connection functions $(\kappa_n)_{n\in\mathbb{N}}$ and some $d \in \mathbb{N}$, assume that there exist $N_0 \in \mathbb{N}$, $t_0 > 0$, $\alpha > d$, and some constant*

$C \in \mathbb{R}^+$ *such that*

$$\sup_{n \geq N_0} \overline{\kappa}_n(t) \leq Ct^{-\alpha} \quad \textit{for all } t \geq t_0,$$

Indeed,

$$\int_0^\infty \phi(t)\, dV_x(t) \leq \int_{t_0}^\infty t^{-\alpha}\, dV_x(t) \leq \int_{t_0}^\infty t^{-\alpha}\, d(ct^d) = cd \int_{t_0}^\infty t^{d-\alpha-1}\, dt, \quad (39)$$

and with $\alpha > d$, this integral converges, satisfying Assumption 7.13. In our theory, we still require the uniform tail control, whereas [18] phrases this condition in terms of the limiting connection function only.

We give a few examples of such polynomial-decay connection functions below in Table 1, with associated identically factorable marking laws given in Table 2. These tables are intended primarily as a catalogue of admissible model families, rather than as a self-contained narrative development. More precisely, they should be read as class-indexed specification tables. They have a common class label $\mathfrak{C}$ which plays the role of a shared identifier, in an almost SQL-style sense, and a valid model specification is obtained by matching the rows with the same class label across the two tables. Each row then represents not a single model, but the family of all instances obtained by instantiating the displayed formulas with parameters and constants that satisfy appropriate constraints, which are detailed in Appendix B.2.

Accordingly, for a fixed class $\mathfrak{C}$, the notation $\mathfrak{C}[\kappa_n(d)]$, $\mathfrak{C}[\kappa_\infty(d)]$, $\mathfrak{C}[\nu]$, and $\mathfrak{C}[*]$ refers to the entries in the corresponding columns of the $\mathfrak{C}$-row, while parentheses in a column heading indicate that the entry is parameterized by the displayed quantity and must be instantiated accordingly. These can be treated as sets of admissible choices, and hence notation such as

$$\kappa_\infty \in \mathfrak{C}[\kappa_\infty(d)]$$

can be used to denote instantiating such a family under suitable constants.

In practice, for a single model of interest, it is sometimes more efficient to verify the assumptions of Proposition 7.15 directly than to work through the full catalogue below. Even when drawing a specification from the tables below, it remains wise to briefly sanity-check that the chosen model is consistent with Proposition 7.15 and Assumption 7.16, since the conditions are interdependent and a small requirement can easily be overlooked.

First, we give the connection function table, which is phrased pre-truncation to $[0, 1]$ for simplicity. Accordingly, each row should be read as implicitly including truncation to $[0, 1]$, even though this is not displayed explicitly in the table itself due to space limitations. The columns $\kappa_n(d)$ and $\kappa_\infty(d)$ are to be read as parameterized row entries. After fixing a class $\mathfrak{C}$, one selects the corresponding $\mathfrak{C}[\kappa_n(d)]$ and $\mathfrak{C}[\kappa_\infty(d)]$ entries and then instantiates them with a common

admissible choice of d and the associated constants from Appendix B.2. The admissible domains of d for each model are shown in the column $\mathfrak{C}[*]$.

When $\cdots$ appears in the $\kappa_n(d)$ column, it indicates that no canonical approximating sequence is prescribed for that class. In that case, one may use the limiting connection function throughout the sequence. If one instead wishes to construct a non-identical approximating sequence for such a row, we leave verification of the assumptions of Proposition 7.15 to the reader.

Class $\mathfrak{C}$	$\kappa_n(d)$	$\kappa_\infty(d)$	$*$
PSIRG	$\cdots$	$f(t)g(w_1,w_2)$	$\mathbb{N}$
GIRG_1	$\cdots$	$\mathbf{1}\left\{\left(\frac{w_1w_2}{\mathbb{E}[M]}\right)^{1/d} > t\right\}$	$\mathbb{N}$
GIRG_2	$\cdots$	$\left(\frac{w_1w_2}{\mathbb{E}[M]}\right)^{\beta} t^{-d\beta}$	$\mathbb{N}$
THRG	$\mathbf{1}\{D_n(t,w_1,w_2) \le R_n\}$	$\mathbf{1}\left\{t \le \frac{\beta w_1 w_2}{\pi}\right\}$	$\{1\}$
PHRG	$\left(1+\exp\left(\frac{D_n(t,w_1,w_2)-R_n}{2T}\right)\right)^{-1}$	$\left(1+\left(\frac{\pi t}{(\beta w_1 w_2)}\right)^{1/T}\right)^{-1}$	$\{1\}$
CSFP	$\cdots$	$1-\exp(-\lambda w_1 w_2 t^{-\alpha})$	$\mathbb{N}$
WDRCM	$\cdots$	$\rho(h(w_1,w_2,t))$	$\mathbb{N}$
SDPARG	$\cdots$	$c(t)t^{-\alpha}$	$\mathbb{N}$

Table 1: Polynomial-distance-scale connection-function families indexed by class $\mathfrak{C}$, with parameterized row entries $\kappa_n(d)$, $\kappa_\infty(d)$, admissible dimension set in column $*$, and shared constant constraints recorded in Appendix B.2.

This table is parameterized by d in the sense that the admissible ranges of the defining constants depend on the ambient dimension, and these restrictions are too detailed to record directly in the table. They are therefore specified separately in Appendix B.2. The admissible domain of d for each class is indicated in the $*$ column, so a row is selected by first fixing a class $\mathfrak{C}$ and then requiring $d \in \mathfrak{C}[*]$. One then instantiates $\mathfrak{C}[\kappa_n(d)]$ and $\mathfrak{C}[\kappa_\infty(d)]$ using constants compatible with that same choice of d. In particular, care is needed in the THRG and PHRG classes, whose admissible dimension set is $\{1\}$.

We do not give a general mechanism for ensuring that κ_∞ is suitably continuous here, as it is difficult to integrate with the table setup, and thus we expect that one can verify it for their choice directly. In the $\mathbb{R}^d$ case, Assumption 7.3 is satisfied whenever κ_∞ is $\lambda_{\mathbb{R}} \otimes \nu \otimes \nu$-a.s. continuous. Indeed, with

$$\Gamma_\infty \overset{d}{=} \delta_x + \eta, \qquad \eta \sim \mathrm{PPP}(c\lambda_{\mathbb{R}^d}),$$

and $A \subseteq \mathbb{R}^+$ is Borel with $\lambda_{\mathbb{R}}(A) = 0$, then

$$D(\Gamma_\infty)(A) = \sum_{p,q \in \Gamma_\infty}^{p \neq q} \mathbf{1}_A(\|p - q\|) \tag{40}$$

$$= 2 \sum_{y \in \eta} \mathbf{1}_A(\|y - x\|) + \sum_{u,v \in \eta}^{u \neq v} \mathbf{1}_A(\|u - v\|). \tag{41}$$

Applying Campbell-Mecke [25, Thm. 4.1, Eq. (4.2)] to the first sum, and twice to the second, gives

$$\begin{aligned}\mathbb{E}[D(\Gamma_\infty)(A)] &= 2\mathbb{E}\left[\sum_{y \in \eta} \mathbf{1}_A(\|y - x\|)\right] + \mathbb{E}\left[\sum_{u,v \in \eta}^{u \neq v} \mathbf{1}_A(\|u - v\|)\right] \\ &= 2c \int_{\mathbb{R}^d} \mathbf{1}_A(\|y - x\|)\, dy + c^2 \iint_{\mathbb{R}^d \times \mathbb{R}^d} \mathbf{1}_A(\|u - v\|)\, du\, dv.\end{aligned}$$

The first term is zero because $\{y \in \mathbb{R}^d : \|y - x\| \in A\}$ has Lebesgue measure zero, and for each fixed $u \in \mathbb{R}^d$ the set $\{v \in \mathbb{R}^d : \|u - v\| \in A\}$ also has Lebesgue measure zero, and thus the second term vanishes. Hence, $\mathbb{E}[D(\Gamma_\infty)(A)] = 0$, and therefore $D(\Gamma_\infty)(A) = 0$ almost surely. Thus, $\lambda_{\mathbb{R}} \otimes \nu \otimes \nu$-a.s. continuity implies $(D(\Gamma_\infty) \otimes \nu \otimes \nu)$-a.s. continuity.

See below the table of corresponding identically factorable mark distributions for the entries in Table 1. This second table is to be read against the first using the same class identifier $\mathfrak{C}$. After fixing a class, one should pair the row labeled $\mathfrak{C}$ here with the row labeled $\mathfrak{C}$ in Table 1, so that the model specification is assembled rowwise across the two tables. Thus, for example, a PSIRG specification is formed from the PSIRG row of Table 1 together with the PSIRG row below, and not by mixing that row with a GIRG, THRG, or other class. In the notation used in the table, this means that $\mathfrak{C}[\nu]$ is selected from the same class $\mathfrak{C}$ for which $\mathfrak{C}[\kappa_n(d)]$ and $\mathfrak{C}[\kappa_\infty(d)]$ were selected.

An entry of $\cdots$ in the mark-distribution column indicates that the construction places no additional restriction on the choice of mark law, while an entry of $\cdots$ in the domain column indicates that no further domain restriction is imposed there. Even in those cases, however, these rows must still be instantiated jointly with Table 1 so as to satisfy the shared constraints recorded in Appendix B.2.

Class $\mathfrak{C}$	ν	**Domain**
PSIRG	$\cdots$	$\cdots$
GIRG_1	$\nu([1, w]) = 1 - w^{1-\beta}$	$[1, \infty)$
GIRG_2	$\nu([1, w]) = 1 - w^{1-\beta}$	$[1, \infty)$
THRG	$\nu([1, w]) = 1 - w^{-2\alpha_H}$	$[1, \infty)$
PHRG	$\nu([1, w]) = 1 - w^{-2\alpha_H}$	$[1, \infty)$

Class $\mathfrak{C}$	ν	**Domain**
CSFP	$\nu([1,w]) = 1 - w^{-\beta}$	$[1,\infty)$
WDRCM	$\nu([0,w]) = w$	$[0,1]$
SDPARG	$\cdots$	$\cdots$

Table 2: Identically factorable mark-law entries ν indexed by the same class key $\mathfrak{C}$ as Table 1, to be matched rowwise with that table under the shared constant constraints recorded in Appendix B.2.

By and large, these connection functions and mark distributions have been adapted from the examples in [18], with the exception of SDPARG, which has a specific use case in the p-adic spaces, covered in Section 7.3.3. The following lemma should therefore be read as a verification of these class-indexed row specifications, instantiated jointly as described above.

▶ **Lemma 7.17.** *Fix $d \in \mathbb{N}$. Choose a model class $\mathfrak{C}$ with $d \in \mathfrak{C}[*]$ and instantiate jointly*

1. *$(\kappa_n)_{n\in\mathbb{N}} \in \mathfrak{C}[\kappa_n(d)]$ via Table 1,*
2. *$(\kappa_\infty) \in \mathfrak{C}[\kappa_\infty(d)]$ via Table 1,*
3. *$(\nu) \in \mathfrak{C}[\nu]$ via Table 2,*

with constants satisfying the shared constraints in Appendix B.2. Then, these choices satisfy Assumptions 7.2 and 7.16.

Proof. The proof follows from the arguments given for each class in Appendix B.2. ◀

As for the actual count sequences that we can handle, our limit is forced to be a Poisson process with some scaled intensity $c\lambda_{\mathbb{R}^d}$, so our model is limited to capturing the extent to which vertex counts within the sequence can handle constrained dispersion. For example, we can handle the following table of count sequences, identified by a class $\mathfrak{D}$.

Class $\mathfrak{D}$	$N(c,(m_n)_{n\in\mathbb{N}})$	$\lambda(c,(m_n)_{n\in\mathbb{N}})$
Fixed	$N_n \overset{d}{=} \lfloor cm_n \rfloor$	$c\lambda_{(S,H)}$
Poisson	$N_n \sim \text{Poisson}(\lfloor cm_n \rfloor)$	$c\lambda_{(S,H)}$
Noisy Binomial	$N_n \sim$ $\text{Bin}(\lceil (1+\varepsilon_n)cm_n \rceil, (1+\varepsilon_n)^{-1})$	$c\lambda_{(S,H)}$
Sublinear Noise	$N_n \overset{d}{=} \lfloor cm_n + cm_n^\alpha Z_n \rfloor$	$c\lambda_{(S,H)}$

Table 3: Concentrating counts and limits for Mixed Binomial vertex processes, parameterized by the sequence $(m_n)_{n\in\mathbb{N}}$ and $c \in \mathbb{R}^+$. The constants are restricted to $(\varepsilon_n)_{n\in\mathbb{N}} \geq 0$ with $\varepsilon_n \to 0$, $0 \leq \alpha < 1$, and Z_n a random non-negative real number with $\sup_n \mathbb{E}[Z_n^3] < \infty$.

Indeed, see the following lemma confirming this. Here the added condition $m_n \to \infty$ simply ensures that the limiting space is not compact, which is needed for the infinite limit. This holds in the $\mathbb{R}^d$ case trivially.

▶ **Lemma 7.18.** *Consider an instance of Context 1. Fix $c \in \mathbb{R}^+$ and let $m_n := \lambda_{(S,H)}(S_n)$ for all $n \in \mathbb{N}$. Choose a count class $\mathfrak{D}$ and instantiate*

$$(N_n)_{n\in\mathbb{N}} \in \mathfrak{D}[N(c, (m_n)_{n\in\mathbb{N}})]$$

with appropriate constants as detailed in caption of Table 3. Then, these choices satisfy Assumption 7.10 whenever $m_n \to \infty$.

The proof is not particularly inspiring and is deferred to Appendix A.8. These tables give us a modular way to construct suitable models. Specifically, in $\mathbb{R}^d$, we have the following corollary of Proposition 7.15.

▶ **Corollary 7.19.** *Fix $c \in \mathbb{R}^+$, $d \in \mathbb{N}$, and $m_n = n$ for all $n \in \mathbb{N}$. Choose a count class $\mathfrak{D}$ and model class $\mathfrak{C}$ with $d \in \mathfrak{C}[*]$. Instantiate jointly*

1. *$(\kappa_n)_{n\in\mathbb{N}} \in \mathfrak{C}[\kappa_n(d)]$ via Table 1,*
2. *$(\kappa_\infty) \in \mathfrak{C}[\kappa_\infty(d)]$ via Table 1, chosen to be $\lambda_{\mathbb{R}} \otimes \nu \otimes \nu$-a.s. continuous,*
3. *$(\nu) \in \mathfrak{C}[\nu]$ via Table 2,*
4. *$(N_n)_{n\in\mathbb{N}} \in \mathfrak{D}[N(c, (m_n)_{n\in\mathbb{N}})]$ via Table 3,*

with constants satisfying the shared constraints in Appendix B.2 and the caption of Table 3. Then, let[17]

$$(\mathbb{R}^{2d}, (\mathbb{T}_n^d)^2_{n\in\mathbb{N}}, 0, \nu, (\kappa_n)_{n\in\mathbb{N}}, \kappa_\infty, (\Gamma_n)_{n\in\mathbb{N}}, \Gamma_\infty, (\xi_n)_{n\in\mathbb{N}}, \xi_\infty)$$

be an instance of Context 1 specializing $(\Gamma_n)_{n\in\mathbb{N}}, \Gamma_\infty$ to

$$\Gamma_n \sim \mathrm{MB}(N_n, \mu_n) \qquad \text{for all } n \in \mathbb{N},$$

$$\Gamma_\infty \overset{d}{=} \delta_x + \eta, \qquad \eta \sim \mathrm{PPP}(c\lambda_{\mathbb{R}^d}),$$

where $(N_n)_{n\in\mathbb{N}}$ is a sequence of random natural numbers and $(\mu_n)_{n\in\mathbb{N}}$ is a sequence of measures given by

$$\mu_n = \frac{1}{n}\lambda_{\mathbb{R}^d}|_{\mathbb{T}_n^d}.$$

Then it holds that

$$G(\xi_n, U_n) \Rightarrow G(\xi_\infty, 0), \qquad U_n \sim \mathbb{U}(\xi_n, \cdot).$$

Proof. We use Proposition 7.15, and thus it suffices to verify Assumptions 7.2, 7.3, 7.7, 7.10 and 7.13. We do each below one by one.

[17] The spaces are given as their square grouped versions. Specifically note that the pointwise torus action on the square space is different from the action of the square torus on itself.

Assumption 7.2: Since we chose $(\kappa_n)_{n\in\mathbb{N}}, \kappa_\infty, \nu$ appropriately from Tables 1 and 2, Lemma 7.17 gives this.

Assumption 7.3: This is given by the assumption on the choice of κ_∞ along with the discussion around Equation (40).

Assumption 7.7: Indeed,

$$[-\frac{n^{1/d}}{2}, \frac{n^{1/d}}{2}] \to \mathbb{R}$$

and thus our subspaces exhaust the ambient space.

Assumption 7.10: Since we chose $(N_n)_{n\in\mathbb{N}}$ appropriately from Table 3, and since $m_n = n = \lambda_{(S,H)}(\mathbb{T}_n^d)$ by construction, Lemma 7.18 gives this.

Assumption 7.13: Since we chose $(\kappa_n)_{n\in\mathbb{N}}, \kappa_\infty, \nu$ appropriately from Tables 1 and 2, Lemma 7.17 gives Assumption 7.16 which then gives this via the surrounding discussion.

The result is then concluded via Proposition 7.15. ◀

7.3.3 Poisson Limit Examples in $\mathbb{Q}_p^d$

Our second example of applicable spaces are the p-adic fields. Fixing some prime p, a p-adic field, $\mathbb{Q}_p$, is a completion of the rationals $\mathbb{Q}$ under a metric known as the p-adic metric. This metric is constructed from a norm via

$$d_p(a,b) = |a-b|_p,$$

where $|\cdot|_p$ relies on the fact that every rational $q \in \mathbb{Q}$ can be expressed as

$$q = p^k \frac{a}{b},$$

where $k \in \mathbb{Z}$ and p does not divide a or b. Under this representation, the norm is given by $|q|_p = p^{-k}$, and is unique. In short, two rationals are "close" in this sense whenever their difference decomposes with a high k in the above decomposition.

Completing $\mathbb{Q}$ under the metric d_p gives the p-adic field $\mathbb{Q}_p$, which is analogous to $\mathbb{R}$ and is LC Polish. The interesting thing about this field is that the closed unit ball is equivalent to the completion of the integers $\mathbb{Z}$ under the same metric, $\mathbb{Z}_p$,

$$\overline{B_{\mathbb{Q}_p}(0,1)} = \mathbb{Z}_p,$$

which is a grouped compact Polish space acting on itself via translation. There is no need for changing the metric like in the torus case, since the sum of any two p-adic points x, y with $|x|_p, |y|_p \leq p^k$ also has $|x+y|_p \leq p^k$ due to the metric being an ultrametric.

A similar property holds for the balls at radius p^n, where $n \in \mathbb{N}$. Specifically

$$\overline{B_{\mathbb{Q}_p}(0, p^n)} = p^{-n}\mathbb{Z}_p,$$

which are also grouped compact Polish spaces. Under the standard normalized Haar measure on the grouped LC Polish space $(\mathbb{Q}_p, +)$, normalized to have

$$\lambda_{\mathbb{Q}_p}(\mathbb{Z}_p) = 1,$$

it holds that

$$\lambda_{\mathbb{Q}_p}(p^{-n}\mathbb{Z}_p) = p^n,$$

for all $n \in \mathbb{N}$. Then $\lambda_{\mathbb{Q}_p^d}$ is simply $\lambda_{\mathbb{Q}_p}^{\otimes d}$. We also extend it to subspaces via restriction, just as in the torus case. Naturally, this is thus also a space where balls of radius r grow polynomially in r, and Assumption 7.16 suffices. Thus, if we consider the β-admissible sequence of subspaces

$$(p^{-n}\mathbb{Z}_p)_{n\in\mathbb{N}}^d \subset_{\beta,x} \mathbb{Q}_p^d,$$

we obtain a result under all of the same count sequences and model examples as the torus case. However, p-adic spaces are usually used for explicitly hierarchical interaction models rather than Euclidean geometry. A closely related benchmark family is ultrametric long-range random graphs where edge probabilities depend only on hierarchical distance, e.g. [8]. While originally studied in the context of percolation, we can use a similar distance-based rule on finite p-adic balls $S_n = (p^{-n}\mathbb{Z}_p)^d$ to arrive at local weak limits, with vertices sampled independently and uniformly in S_n (under any of the count regimes above). The corresponding connection function κ can be written as

$$\kappa(p^k) := \min\left\{ \frac{c_k}{p^{dk(1+\delta)}}, 1 \right\}, \quad k \in \mathbb{Z},$$

where $\delta > 0$ and $(c_k)_{k\in\mathbb{Z}}$ is a chosen positive sequence controlling the connection strength at each distance shell k (for example, $c_k \equiv c$ gives a pure power-law tail in k). Any continuous interpolation for values of t not of this form is sufficient to apply the model, as it is irrelevant for the p-adic metric, and thus we can assume κ is continuous on $\mathbb{R}^+$ without loss of generality.

If $(c_k)_{k\in\mathbb{Z}}$ is bounded via $\sup_{k\in\mathbb{Z}} c_k \le C$, then for all sufficiently large $t \in p^{\mathbb{Z}}$,

$$\kappa(t) \le C\, t^{-d(1+\delta)},$$

so the model stays in the same power-law decay class. Also, on $\mathbb{Q}_p^d$, the ball-volume profile changes only when the radius reaches a power of p, so the associated Stieltjes measure is concentrated on those shell radii. Consequently, the decay condition is determined entirely by shell values at powers of p, and any interpolation between shells does not affect it. In this form, it is exactly a p-adic SIRG and Proposition 7.15 applies. We refer to such a model as a shell-decay

p-adic random graph, or SDPARG for short, which is equivalent to the form in Table 1.

This p-adic model is markless, and can easily be extended with marks under our framework. As a quick hypothetical example, one can attach a node reliability mark and let high-reliability pairs connect with slightly higher probability at every shell. In the case of future extension to more marking kernel classes, rather than just identically factorable marking kernels, these types of extensions become even more viable.

One may also use any other connection function from Table 1 in the p-adic setting, though the utility of doing so is not detailed here. Since the metric takes on values only in $p^{\mathbb{Z}} = \{p^z \mid z \in \mathbb{Z}\}$, continuity on this discrete set suffices. Thus, we have an analogous corollary in the p-adic setting:

▶ **Corollary 7.20.** *Fix $c \in \mathbb{R}^+$, $d \in \mathbb{N}$, and $m_n = p^{nd}$ for all $n \in \mathbb{N}$. Choose a count class $\mathfrak{D}$ and model class $\mathfrak{C}$ with $d \in \mathfrak{C}[*]$. Instantiate jointly*

1. *$(\kappa_n)_{n\in\mathbb{N}} \in \mathfrak{C}[\kappa_n(d)]$ via Table 1,*
2. *$(\kappa_\infty) \in \mathfrak{C}[\kappa_\infty(d)]$ via Table 1, chosen to be $p^{\mathbb{Z}} \otimes \nu \otimes \nu$-a.s. continuous,*
3. *$(\nu) \in \mathfrak{C}[\nu]$ via Table 2,*
4. *$(N_n)_{n\in\mathbb{N}} \in \mathfrak{D}[N(c, (m_n)_{n\in\mathbb{N}})]$ via Table 3,*

with constants satisfying the shared constraints in Appendix B.2 and the caption of Table 3. Then, let

$$(\mathbb{Q}_p^{2d}, ((p^{-n}\mathbb{Z}_p)^d)^2_{n\in\mathbb{N}}, 0, \nu, (\kappa_n)_{n\in\mathbb{N}}, \kappa_\infty, (\Gamma_n)_{n\in\mathbb{N}}, \Gamma_\infty, (\xi_n)_{n\in\mathbb{N}}, \xi_\infty)$$

be an instance of Context 1 specializing $(\Gamma_n)_{n\in\mathbb{N}}, \Gamma_\infty$ to

$$\Gamma_n \sim \mathrm{MB}(N_n, \mu_n) \qquad \text{for all } n \in \mathbb{N},$$

$$\Gamma_\infty \overset{d}{=} \delta_x + \eta, \qquad \eta \sim \mathrm{PPP}(c\lambda_{\mathbb{Q}_p^d}),$$

where $(N_n)_{n\in\mathbb{N}}$ is a sequence of random natural numbers and $(\mu_n)_{n\in\mathbb{N}}$ is a sequence of measures given by

$$\mu_n = \frac{1}{p^{nd}} \lambda_{\mathbb{Q}_p^d}|_{(p^{-n}\mathbb{Z}_p)^d}.$$

Then it holds that

$$G(\xi_n, U_n) \Rightarrow G(\xi_\infty, 0), \qquad U_n \sim \mathbb{U}(\xi_n, \cdot).$$

Proof. We use Proposition 7.15, and thus it suffices to verify Assumptions 7.2, 7.3, 7.7, 7.10 and 7.13. We do each below one by one.

Assumption 7.2: Since we chose $(\kappa_n)_{n\in\mathbb{N}}, \kappa_\infty, \nu$ appropriately from Tables 1 and 2, Lemma 7.17 gives this.

Assumption 7.3: This is given by the assumption on the choice of κ_∞, since on $\mathbb{Q}_p^d$ the distance coordinate takes values only in $p^{\mathbb{Z}}$.

Assumption 7.7: Indeed,

$$p^{-n}\mathbb{Z}_p \to \mathbb{Q}_p$$

and thus our subspaces exhaust the ambient space.

Assumption 7.10: Since we chose $(N_n)_{n\in\mathbb{N}}$ appropriately from Table 3, and since $m_n = p^{nd} = \lambda_{(S,H)}((p^{-n}\mathbb{Z}_p)^d)$ by construction, Lemma 7.18 gives this.

Assumption 7.13: Since we chose $(\kappa_n)_{n\in\mathbb{N}}, \kappa_\infty, \nu$ appropriately from Tables 1 and 2, and since $\mathbb{Q}_p^d$ also exhibits polynomial growth, Assumption 7.16 is sufficient and is given by Lemma 7.17.

Thus, Proposition 7.15 gives the result. ◀

7.4 Cox Vertex Processes with Cox Limits

7.4.1 General Theory for Cox Limits

In this paradigm, we consider instead a sequence of Cox vertex processes converging to a limiting Cox vertex process. This is what allows us to finally have Cox limiting objects, which is very useful as random graphs constructed on Cox vertex processes are highly general. For motivation, fix for each $n \in \mathbb{N}$,

$$\Gamma_n \sim \mathrm{Cox}(\Phi_n),$$

for some directing measure Φ_n. This sequence naturally has a Cox limit by [19, Thm. 23.21]. More specifically, assuming that $\Phi_n \Rightarrow \Psi$ for some limiting directing measure, it follows that

$$\Gamma_n \Rightarrow \eta, \qquad \eta \sim \mathrm{Cox}(\Psi).$$

Naturally, this extends analogously to the Palm variants via Equation (34). Specifically, for some $x \in S$, if Φ_n^x exists uniquely for each $n \in \mathbb{N}$ and $\Phi_n^x \Rightarrow \Phi_\infty$ for some Φ_∞, then

$$\Gamma_n^x \Rightarrow \delta_x + \eta, \qquad \eta \sim \mathrm{Cox}(\Phi_\infty).$$

In the Euclidean case used below, stationarity of the directing measure implies stationarity of the Cox process by [32, p. 725]. Hence, the following assumption and lemma are sufficient to conclude Assumption 7.1. Note that just like in our discussion around Definition 20, we treat directing measures as existing canonically on both their subspace and the ambient space, so we can discuss both stationarity and convergence. In this case, we extend them to the ambient space via zero extension, just like with point processes.

▶ **Assumption 7.21.** *Given an instance of Context 1 and a sequence $(\Phi_n)_{n\in\mathbb{N}}$ of directing measures, each on (S_n, H_n), and some directing measure Φ_∞ on (S, H), assume that all Φ_n are stationary on their respective spaces and that $\Phi_n^x \Rightarrow \Phi_\infty$.*

▶ **Lemma 7.22.** *Let the tuple*

$$((S,H)^2,(S_n,H_n)^2_{n\in\mathbb{N}},x,\nu,(\kappa_n)_{n\in\mathbb{N}},\kappa_\infty,(\Gamma_n)_{n\in\mathbb{N}},\Gamma_\infty,(\xi_n)_{n\in\mathbb{N}},\xi_\infty)$$

be an instance of Context 1 specializing $(\Gamma_n)_{n\in\mathbb{N}},\Gamma_\infty$ *to*

$$\Gamma_n \sim \mathrm{Cox}(\Phi_n) \qquad \text{for all } n\in\mathbb{N},$$

$$\Gamma_\infty \overset{d}{=} \delta_x+\eta, \qquad \eta\sim\mathrm{Cox}(\Phi_\infty),$$

where $(\Phi_n)_{n\in\mathbb{N}},\Phi_\infty$ *are directing measures. Then, Assumption 7.21 gives Assumption 7.1.*

Proof. By Assumption 7.21 and [32, p.725], each $\Gamma_n\sim\mathrm{Cox}(\Phi_n)$ is stationary on (S_n,H_n). For the Palm convergence, Equation (34) gives

$$\Gamma_n^x \overset{d}{=} \delta_x+\eta_n, \qquad \eta_n\sim\mathrm{Cox}(\Phi_n^x).$$

By Assumption 7.21, we have $\Phi_n^x\Rightarrow\Phi_\infty$, and thus [19, Thm. 23.21] yields $\eta_n\Rightarrow\eta$, where $\eta\sim\mathrm{Cox}(\Phi_\infty)$. Therefore,

$$\Gamma_n^x\Rightarrow\delta_x+\eta=\Gamma_\infty.$$

Hence, Assumption 7.21 gives Assumption 7.1. ◀

One may notice that Assumption 7.7 is no longer needed here. This is because the convergence of the Palm directing measures $\Phi_n^x\Rightarrow\Phi_\infty$ is already assumed directly, so the ambient limiting object is built into the hypothesis rather than recovered from the subspaces, as in the mixed binomial case. Also, this assumption does not require the directing measures to be diffuse, while we will keep this requirement for the later results, to ensure simplicity of the resulting vertex processes.

Unfortunately, these directing measures may be highly general, and as such most of the other conditions, such as vertex concentration and continuity assumptions, cannot be easily reduced. Fortunately though, Assumption 7.13 still can be used here to resolve the challenging radial tightness and fiber-finiteness conditions. This requires one extra moment assumption to go through, given in a way that is different than the mixed binomial case. Intuitively, we want to ensure that the intensity and higher factorial moments are dominated by a constant factor of the ambient measure, which allows us to reduce the integrals of Assumption 7.5 to integrals on the ambient space. It is plausible that a more intricate proof could reduce these conditions in future work. See below the assumption and lemma.

▶ **Assumption 7.23.** *Given a sequence* $(\Phi_n)_{n\in\mathbb{N}}$ *of directing measures all*[18] *on some grouped LC Polish space* (S,H)*, and some* $x\in S$*, assume that there exist*

[18] This is an equivalent view to subspaces via zero extension.

constants $c_1, c_2, c_3 > 0$ and $N_0 \in \mathbb{N}$ such that for all $n \geq N_0$ and all relatively compact Borel sets $A, B, C \subseteq S$,

$$\begin{aligned}\mathbb{E}[\Phi_n^x(A)] &\leq c_1\, \lambda_{(S,H)}(A),\\ \mathbb{E}[\Phi_n^x(A)\Phi_n^x(B)] &\leq c_2\, \lambda_{(S,H)}(A)\lambda_{(S,H)}(B),\\ \mathbb{E}[\Phi_n^x(A)\Phi_n^x(B)\Phi_n^x(C)] &\leq c_3\, \lambda_{(S,H)}(A)\lambda_{(S,H)}(B)\lambda_{(S,H)}(C).\end{aligned}$$

whenever these Palm versions exist uniquely.

▶ **Lemma 7.24.** *Let the tuple*

$$((S,H)^2, (S_n,H_n)^2_{n\in\mathbb{N}}, x, \nu, (\kappa_n)_{n\in\mathbb{N}}, \kappa_\infty, (\Gamma_n)_{n\in\mathbb{N}}, \Gamma_\infty, (\xi_n)_{n\in\mathbb{N}}, \xi_\infty)$$

be an instance of Context 1 specializing $(\Gamma_n)_{n\in\mathbb{N}}, \Gamma_\infty$ to

$$\Gamma_n \sim \mathrm{Cox}(\Phi_n) \qquad \text{for all } n \in \mathbb{N},$$

$$\Gamma_\infty \overset{d}{=} \delta_x + \eta, \qquad \eta \sim \mathrm{Cox}(\Phi_\infty),$$

where $(\Phi_n)_{n\in\mathbb{N}}, \Phi_\infty$ are diffuse directing measures. Then, Assumptions 7.13, 7.21 and 7.23 together give Assumption 7.5.

The proof is analogous to Lemma 7.14 and is deferred to Appendix A.7. With this, we give the proposition for Cox vertex processes in the most general form within the scope of this thesis.

▶ **Proposition 7.25.** *Let the tuple*

$$((S,H)^2, (S_n,H_n)^2_{n\in\mathbb{N}}, x, \nu, (\kappa_n)_{n\in\mathbb{N}}, \kappa_\infty, (\Gamma_n)_{n\in\mathbb{N}}, \Gamma_\infty, (\xi_n)_{n\in\mathbb{N}}, \xi_\infty)$$

be an instance of Context 1 specializing $(\Gamma_n)_{n\in\mathbb{N}}, \Gamma_\infty$ to

$$\Gamma_n \sim \mathrm{Cox}(\Phi_n) \qquad \text{for all } n \in \mathbb{N},$$

$$\Gamma_\infty \overset{d}{=} \delta_x + \eta, \qquad \eta \sim \mathrm{Cox}(\Phi_\infty),$$

where $(\Phi_n)_{n\in\mathbb{N}}, \Phi_\infty$ are diffuse directing measures. Then, under Assumptions 7.2–7.4, 7.13, 7.21 and 7.23, it holds that

$$G(\xi_n, U_n) \Rightarrow G(\xi_\infty, x), \qquad U_n \sim \mathbb{U}(\xi_n, \cdot).$$

Proof. We use Corollary 7.6. Thus, we need to verify Assumptions 7.1–7.5. We do this one by one below:

Assumption 7.1: This is given by Lemma 7.22 under Assumption 7.21.

Assumption 7.2 This remains assumed directly.

Assumption 7.3 This remains assumed directly.

Assumption 7.4: This remains assumed directly.

Assumption 7.5: This is given by Lemma 7.24 under Assumptions 7.13, 7.21 and 7.23.

This concludes the proof via application of Corollary 7.6. ◀

For the Cox to Cox paradigm, we will only give one example in $\mathbb{R}^d$, as an illustrative demonstration of the results. This is not meant to be exhaustive of the framework's capacity. We expect that the exposition provided throughout this section should be sufficient to assist in one's efforts to extend these Cox results to p-adic spaces or more exotic directing measures in future work.

7.4.2 Cox Limit Example in $\mathbb{R}^d$

We again consider the torus setup from earlier, although in this case we will be interested in the n^d volume tori $\mathbb{T}^d_{n^d}$ for reasons that will be mentioned shortly. We will demonstrate the Cox framework via one of the most elementary classes of stationary directing measures, which we refer to a *phase-shifted* directing measures. Formally, for each $n \in \mathbb{N}$, we choose some measurable $q_n : \mathbb{R}^d \to [0, \infty)$, satisfying the following two properties:

▶ **Assumption 7.26.** *Given a sequence of measurable functions $(q_n)_{n\in\mathbb{N}} : \mathbb{R}^d \to [0, \infty)$, assume that*

1. *For each $n \in \mathbb{N}$,*
$$\int_{[-1/2,1/2)^d} q_n(u)\lambda_{\mathbb{R}^d}(du) = 1,$$
2. *For any $x \in \mathbb{R}^d$, $y \in \mathbb{Z}^d$, it holds for each $n \in \mathbb{N}$ that*
$$q_n(x + y) = q_n(x).$$

The first property is not strictly necessary for the construction that will follow, but is standard as it highlights the distribution of mass within each period. The second property is referred to as $\mathbb{Z}^d$-periodicity, and is effectively what ensures stationarity in the following construction. Intuitively, these *profiles*, as we will call them henceforth, represent a fixed local pattern of mass which is repeated identically across every unit cell of the unit lattice, so that the global environment is built by tiling space with the same microscopic structure.

Some examples of such profiles in $d = 1$ are below. For each of these, normalization against their $[-1/2, 1/2)$ integral is necessary before they can be applied.

Expression	**Variables**
$c_1 + c_2 \sin(2\pi u) + c_3 \cos(4\pi u)$	$c_1 > \lvert c_2\rvert + \lvert c_3\rvert$
$c_1 + c_2 \sin^2(2\pi u) + c_3 \cos^2(4\pi u)$	$c_1 > 0,\ c_2, c_3 \geq 0$
$c_1 + \exp(c_2 \sin(2\pi u) + c_3 \cos(4\pi u))$	$c_1 \geq 0,\ c_2, c_3 \in \mathbb{R}$

Table 4: Non-normalized positive $\mathbb{Z}$-periodic profiles in $d = 1$.

Some examples of such profiles in $d = 2$ are also included. For these, normalization must occur against their $[-1/2, 1/2)^2$ integral instead.

Expression	Variables
$c_1 + c_2 \sin(2\pi u_1) + c_3 \cos(4\pi u_2)$	$c_1 > \lvert c_2\rvert + \lvert c_3\rvert$
$c_1 + c_2 \sin(2\pi u_1)\sin(2\pi u_2) + c_3 \cos(4\pi u_1)\cos(4\pi u_2)$	$c_1 > \lvert c_2\rvert + \lvert c_3\rvert$
$c_1 + c_2 \sin^2(2\pi u_1) + c_3 \cos^2(4\pi u_2)$	$c_1 > 0,\ c_2, c_3 \geq 0$
$c_1 + \exp(c_2 \sin(2\pi u_1) + c_3 \cos(4\pi u_2))$	$c_1 \geq 0,\ c_2, c_3 \in \mathbb{R}$

Table 5: Non-normalized positive $\mathbb{Z}^2$-periodic profiles in $d = 2$.

The proof that these are periodic is left to the reader, as it follows almost immediately from the periodicity of sin and cos.

Once such a sequence of profiles has been chosen, we can take a uniformly distributed shift variable Θ on $[-1/2, 1/2)^d$, and define the directing measure by

$$\Phi_n(A) := \int_A q_n(u + \Theta)\lambda_{\mathbb{T}^d_{n^d}}(du), \tag{42}$$

for every bounded Borel $A \subseteq \mathbb{R}^d$. Note that here we are effectively doing zero extension outside of the tori, since we recall that the torus reference measures are restrictions of the ambient reference measures. Thus, the periodic density pattern is essentially shifted around uniformly, which is exactly what gives stationarity. This also illuminates why we need the n^d volume tori, as otherwise the tori would not naturally rest in the periodic intervals, and stationarity would not follow. To formalize this, we have the following lemma.

▶ **Lemma 7.27.** *Given a sequence of measurable profiles $(q_n)_{n\in\mathbb{N}} : \mathbb{R}^d \to [0, \infty)$ satisfying Assumption 7.26, it holds that the directing measures given by Equation (42) are each stationary on $\mathbb{T}^d_{n^d}$.*

Proof. Fix $n \in \mathbb{N}$ and $h \in \mathbb{T}^d_{n^d}$. It is enough to show that $\theta_h \Phi_n \overset{d}{=} \Phi_n$. For any bounded Borel $A \subseteq \mathbb{R}^d$, translation invariance of $\lambda_{\mathbb{T}^d_{n^d}}$ gives

$$\begin{aligned}\theta_h \Phi_n(A) = \Phi_n(A - h) &= \int_{A-h} q_n(u + \Theta)\lambda_{\mathbb{T}^d_{n^d}}(du) \\ &= \int_A q_n(v - h + \Theta)\lambda_{\mathbb{T}^d_{n^d}}(dv).\end{aligned}$$

Now let $\widetilde{\Theta}$ be the representative of $\Theta - h$ in $[-1/2, 1/2)^d$ modulo $\mathbb{Z}^d$. Since Θ is uniform on $[-1/2, 1/2)^d$, also $\widetilde{\Theta}$ is uniform on $[-1/2, 1/2)^d$, and by $\mathbb{Z}^d$-periodicity of q_n we have

$$q_n(v - h + \Theta) = q_n(v + \widetilde{\Theta})$$

for all $v \in \mathbb{R}^d$. Therefore,

$$\theta_h \Phi_n(A) \overset{d}{=} \int_A q_n(v + \widetilde{\Theta})\lambda_{\mathbb{T}^d_{n^d}}(dv) \overset{d}{=} \int_A q_n(v + \Theta)\lambda_{\mathbb{T}^d_{n^d}}(dv) = \Phi_n(A).$$

Since this holds for every bounded Borel A, it follows that $\theta_h \Phi_n \overset{d}{=} \Phi_n$, so Φ_n is stationary on $\mathbb{T}^d_{n^d}$. ◀

Furthermore, it holds that the Palm variants have a representation in terms of these profiles. See the following lemma.

▶ **Lemma 7.28.** *Let $q : \mathbb{R}^d \to [0,\infty)$ be a measurable profile satisfying*[19] *Assumption 7.26. Then, for Φ constructed as in Equation (42), it holds that for any $n \in \mathbb{N}$,*

$$\Phi^0(A) := \int_A q(u + \Theta^0) \lambda_{\mathbb{T}^d_{n^d}}(du)$$

for all bounded Borel $A \subseteq \mathbb{R}^d$, where Θ^0 satisfies

$$\mathbb{P}(\Theta^0 \in B) = \int_B q(\theta) \lambda_{\mathbb{R}^d}(d\theta)$$

for every Borel $B \subseteq [-1/2, 1/2)^d$.

Proof. Write $\phi_\theta(A) := \int_A q(u+\theta)\lambda_{\mathbb{T}^d_{n^d}}(du)$, so $\Phi = \phi_\Theta$. By Assumption 7.26,

$$\Lambda_\Phi(A) = \mathbb{E}[\Phi(A)] = \int_A \mathbb{E}[q(u+\Theta)]\lambda_{\mathbb{T}^d_{n^d}}(du) = \lambda_{\mathbb{T}^d_{n^d}}(A).$$

Now let Θ^x have density $\theta \mapsto q(x+\theta)$ on $[-1/2,1/2)^d$, and set $\Phi^x := \phi_{\Theta^x}$. Then for every nonnegative measurable $f : S \times \mathcal{M}(S) \to [0,\infty)$,

$$\begin{aligned}\mathbb{E}\left[\int_S f(y,\Phi)\,\Phi(dy)\right] &= \int_S \int_{[-1/2,1/2)^d} f(y,\phi_\theta) q(y+\theta) \lambda_{\mathbb{R}^d}(d\theta)\lambda_{\mathbb{T}^d_{n^d}}(dy) \\ &= \int_S \mathbb{E}[f(y,\Phi^y)]\,\Lambda_\Phi(dy).\end{aligned}$$

So $(\Phi^x)_{x\in S}$ satisfies Equation (33). In particular, Θ^0 has law

$$\mathbb{P}(\Theta^0 \in B) = \int_B q(\theta)\lambda_{\mathbb{R}^d}(d\theta),$$

and the conclusion follows. ◀

Hence, we are just about ready to give our corollary. The last assumption we need is to guarantee convergence and boundedness of our profile sequence. This is what will ultimately give us convergence of the Cox processes and the moment conditions of Assumption 7.23.

▶ **Assumption 7.29.** *Given a sequence of measurable profiles $(q_n)_{n\in\mathbb{N}} : \mathbb{R}^d \to [0,\infty)$ and measurable profile q_∞ all satisfying Assumption 7.26, assume that $q_n \to q_\infty$ uniformly and that*

$$\sup_{n\in\mathbb{N}} \sup_{x\in[0,1)^d} q_n(x) < \infty.$$

[19] Technically this assumption is for sequences but applies identically here.

At last, we give a corollary for these models. As in the mixed-binomial examples, the kernel and mark law are chosen from the tables. There is, however, no separate count-class choice here, since the vertex counts are induced by the Cox directing measures themselves.

▶ **Corollary 7.30.** *Fix $d \in \mathbb{N}$, and choose a model class $\mathfrak{C}$ with $d \in \mathfrak{C}[*]$. Instantiate jointly*

1. *$(\kappa_n)_{n\in\mathbb{N}} \in \mathfrak{C}[\kappa_n(d)]$ via Table 1,*
2. *$\kappa_\infty \in \mathfrak{C}[\kappa_\infty(d)]$ via Table 1, chosen to be $\lambda_{\mathbb{R}} \otimes \nu \otimes \nu$-a.s. continuous,*
3. *$\nu \in \mathfrak{C}[\nu]$ via Table 2,*

with constants satisfying the shared constraints in Appendix B.2. Let $(q_n)_{n\in\mathbb{N}}$ and q_∞ satisfy Assumptions 7.26 and 7.29. For each $n \in \mathbb{N}$, define

$$\Phi_n(A) := \int_A q_n(u + \Theta_n)\lambda_{\mathbb{T}^d_{n^d}}(du),$$

where $\Theta_n \sim \mathrm{Unif}([-1/2, 1/2)^d)$. Define

$$\Phi_\infty(A) := \int_A q_\infty(u + \Theta^0_\infty)\lambda_{\mathbb{R}^d}(du)$$

where Θ^0_∞ has density q_∞ on $[-1/2, 1/2)^d$. Then let

$$(\mathbb{R}^{2d}, (\mathbb{T}^d_{n^d})^2_{n\in\mathbb{N}}, 0, \nu, (\kappa_n)_{n\in\mathbb{N}}, \kappa_\infty, (\Gamma_n)_{n\in\mathbb{N}}, \Gamma_\infty, (\xi_n)_{n\in\mathbb{N}}, \xi_\infty)$$

be an instance of Context 1 specializing $(\Gamma_n)_{n\in\mathbb{N}}, \Gamma_\infty$ to

$$\Gamma_n \sim \mathrm{Cox}(\Phi_n) \qquad \text{for all } n \in \mathbb{N},$$

$$\Gamma_\infty \overset{d}{=} \delta_0 + \eta, \qquad \eta \sim \mathrm{Cox}(\Phi_\infty).$$

Then,

$$G(\xi_n, U_n) \Rightarrow G(\xi_\infty, 0), \qquad U_n \sim \mathbb{U}(\xi_n, \cdot).$$

Proof. We use Proposition 7.25. Thus, it suffices to verify Assumptions 7.2–7.4, 7.13, 7.21 and 7.23. Note that these directing measures are naturally diffuse since they are defined via integration against diffuse measures.

Assumption 7.2: Since we chose $(\kappa_n)_{n\in\mathbb{N}}, \kappa_\infty, \nu$ appropriately from Tables 1 and 2, Lemma 7.17 gives this.

Assumption 7.3: The same argument as in Equation (40) applies here after conditioning on Θ^0_∞, since conditional on Θ^0_∞ the Cox limit is a Poisson point process with directing measure given by the density $q_\infty(\cdot + \Theta^0_\infty)$ against Lebesgue measure. Thus, the assumed $(\lambda_{\mathbb{R}^d} \otimes \nu \otimes \nu)$-almost sure continuity of κ_∞ implies Assumption 7.3.

Assumption 7.13: Since we chose $(\kappa_n)_{n\in\mathbb{N}}, \kappa_\infty, \nu$ appropriately from Tables 1 and 2, Lemma 7.17 gives Assumption 7.16, which implies this assumption by the surrounding discussion.

Assumption 7.23: Let

$$C := \sup_{n\in\mathbb{N}} \sup_{x\in[0,1)^d} q_n(x) < \infty.$$

By $\mathbb{Z}^d$-periodicity, $q_n \le C$ on $\mathbb{R}^d$ for every n. By the previous lemma, for every relatively compact Borel $A \subseteq \mathbb{R}^d$,

$$\Phi_n^0(A) = \int_A q_n(u+\Theta_n^0)\lambda_{\mathbb{T}^d_{n^d}}(du) \le C\lambda_{\mathbb{R}^d}(A).$$

Hence, for relatively compact Borel $A, B, C' \subseteq \mathbb{R}^d$,

$$\begin{aligned}
\mathbb{E}[\Phi_n^0(A)] &\le C\lambda_{\mathbb{R}^d}(A),\\
\mathbb{E}[\Phi_n^0(A)\Phi_n^0(B)] &\le C^2\lambda_{\mathbb{R}^d}(A)\lambda_{\mathbb{R}^d}(B),\\
\mathbb{E}[\Phi_n^0(A)\Phi_n^0(B)\Phi_n^0(C')] &\le C^3\lambda_{\mathbb{R}^d}(A)\lambda_{\mathbb{R}^d}(B)\lambda_{\mathbb{R}^d}(C').
\end{aligned}$$

So Assumption 7.23 holds.

Assumption 7.4: Also,

$$\Phi_n(\mathbb{T}^d_{n^d}) = \int_{\mathbb{T}^d_{n^d}} q_n(u+\Theta_n)\lambda_{\mathbb{T}^d_{n^d}}(du) = n^d$$

almost surely, since $\mathbb{T}^d_{n^d}$ is the union of n^d unit cells and q_n has unit mass on each cell. Therefore

$$|\Gamma_n| \sim \text{Poisson}(n^d),$$

and

$$\mathbb{E}\left[\left|\frac{|\Gamma_n|}{\mathbb{E}[|\Gamma_n|]} - 1\right|\right] \le \frac{\sqrt{\text{Var}(|\Gamma_n|)}}{\mathbb{E}[|\Gamma_n|]} = n^{-d/2} \to 0.$$

Thus Assumption 7.4 holds.

Assumption 7.21: By the stationarity lemma above, each Φ_n is stationary on $\mathbb{T}^d_{n^d}$. By the Palm-profile lemma above,

$$\Phi_n^0(A) = \int_A q_n(u+\Theta_n^0)\lambda_{\mathbb{T}^d_{n^d}}(du),$$

where Θ_n^0 has density q_n on $[-1/2,1/2)^d$. Since $q_n \to q_\infty$ uniformly,

$$\|\mathcal{L}(\Theta_n^0) - \mathcal{L}(\Theta_\infty^0)\|_{\text{TV}} \le \int_{[-1/2,1/2)^d} |q_n(\theta) - q_\infty(\theta)|\, d\theta \to 0,$$

so $\Theta_n^0 \Rightarrow \Theta_\infty^0$. Now fix $f \in C_c(\mathbb{R}^d)$ with compact support K. For all sufficiently large n, $K \subseteq \mathbb{T}^d_{n^d}$. Define

$$\psi_n(\theta) := \int_K f(u)q_n(u+\theta)\, du, \qquad \psi_\infty(\theta) := \int_K f(u)q_\infty(u+\theta)\, du.$$

Then,

$$\sup_{\theta\in[-1/2,1/2)^d} |\psi_n(\theta) - \psi_\infty(\theta)| \leq \|f\|_\infty \lambda_{\mathbb{R}^d}(K) \|q_n - q_\infty\|_\infty \to 0,$$

and ψ_∞ is continuous because translations are continuous. Hence,

$$\int f \, d\Phi_n^0 = \psi_n(\Theta_n^0) \Rightarrow \psi_\infty(\Theta_\infty^0) = \int f \, d\Phi_\infty.$$

Therefore, $\Phi_n^0 \Rightarrow \Phi_\infty$, so Assumption 7.21 holds.

The result now follows from Proposition 7.25. ◀

8 Discussion

With the mathematics now complete and verified via examples, we give a review and discussion. This will cover the results found, the context in which these results can be applied, and possible next steps and future extensions of the work. We begin with the summary of results.

8.1 Summary of Results

The primary goal of this paper was to establish a modular framework for determining weak convergence of a wide range of random rooted graph models. We presented this work in three levels of generality, each becoming more specific and applicable than the last. Perhaps the most central result that we established is Theorem 4.1, our weak convergence theorem, as it formalizes the equivalence between graph convergence and radial tightness. This theorem is central because it isolates a mechanism that appears repeatedly in proofs of local weak convergence for undirected spatial graphs, namely the need to prevent finite rooted neighborhoods from depending on vertices at arbitrarily large spatial distance. In our formulation, this control is expressed through radial tightness. The advantage of Theorem 4.1 is therefore that it separates this reusable convergence mechanism from the model-specific estimates needed to verify it.

Even if a proof were to be carried out directly for a specific model, it would often need to establish a condition playing essentially the same role, and then use a comparable localization argument. For example, a similar no-escape mechanism appears in local limits of spatial Gibbs random graphs. Endo and Valesin prove convergence via truncated balls, with the key long-edge control given by [11, Prop. 2.1, Eq. (2.5)], and then pass to full neighborhoods by sending the cutoff to infinity in the proof of their Theorem 1.2. This is analogous in role to radial tightness, as both arguments prevent rooted neighborhoods from escaping to arbitrarily large spatial scales, and it is probable that the result could be replicated via Theorem 4.1, just as the result of [18] was replicated in Section 7.3. However, we do not formally show this. It is also likely that many combinatorial models (i.e. models that rely on an indexed set of vertices) can also be covered by assigning uniformly selected labels to vertices, though this is not formally verified here.

Thus, Theorem 4.1 provides the very core of the machinery, as any spatial undirected random graph model can be analyzed[20] through it. The results of Section 5, Section 6, and Section 7 therefore mainly serve to reduce the conditions of Theorem 4.1 to easily checkable variants in specific model contexts.

The second result is the local weak convergence framework under uniform rooting developed in Section 5, culminating in Theorem 5.2. Conceptually, this is where the paper turns local weak convergence into a tractable pipeline. Instead

[20] The adjacency measure representation naturally allows for such analysis, but we do not claim that our framework is the most effective route for any given model.

of working directly with a random uniformly chosen vertex in each finite graph, we pass to a fixed-root Palm viewpoint on the corresponding adjacency processes. This reduction matters because Palm distributions isolate the effect of root selection into a structured probabilistic transformation, so the remaining convergence analysis can be carried out at a fixed location. In this way, uniform root selection is no longer an external complication, and instead becomes part of the model via Palm theory. Combined with Theorem 4.1, this yields a modular decomposition of the full task into fixed-root process convergence, regularity/fibre-finiteness at the limit, and radial tightness of neighborhoods.

The third result is the specialization to random graph models in Section 6, where SIRGs are treated as a broad connection-kernel class rather than as isolated examples. The connection-kernel viewpoint is equivalent in reach to the direct adjacency process viewpoint, but significantly better for analysis. A random graph law can thus be represented through a vertex process together with a measurable connection rule, and this representation makes assumptions explicit, modular, and transferable across models.

Within this setup, Theorem 6.4 and its intermediate results, Lemmas 6.6 and 6.7, convert the difficult conditions in the abstract theorem into concrete checks, with fibre-finiteness handled through integrability and moment bounds and radial tightness handled through explicit tail control on long-range edge contributions. This is a practical advance, because these checks are directly verifiable in concrete families and no longer require rebuilding a full local limit proof from scratch for each model.

Building on these conditions, the examples section, Section 7, is not meant to be exhaustive, but its role is rather to demonstrate that the framework scales and to show how existing local limit results can be recovered or generalized through a single method. To this end we provide examples for Poisson and Cox vertex processes, along with examples in p-adic vertex spaces. The main message is that once the framework is in place, proving convergence for a new model class is primarily a matter of checking a finite list of conditions. This shift from ad hoc proof engineering to reusable condition checking is exactly the long-term research practice this paper is intended to establish.

Another substantive point is that the layers are not merely weaker or stronger statements, but different interfaces between probability structure and graph level convergence. The weak convergence theorem is extremely general but demands abstract regularity and tightness checks, while the higher layers trade that generality for assumptions that can be verified through moment and tail estimates. This means that model placement in the pipeline is effectively a choice of proof interface.

The figure below makes this interface structure explicit across the model families considered here. By identifying where a prospective model fits relative to the theory, one can identify which theorem can be applied, and which conditions must be verified for that model.

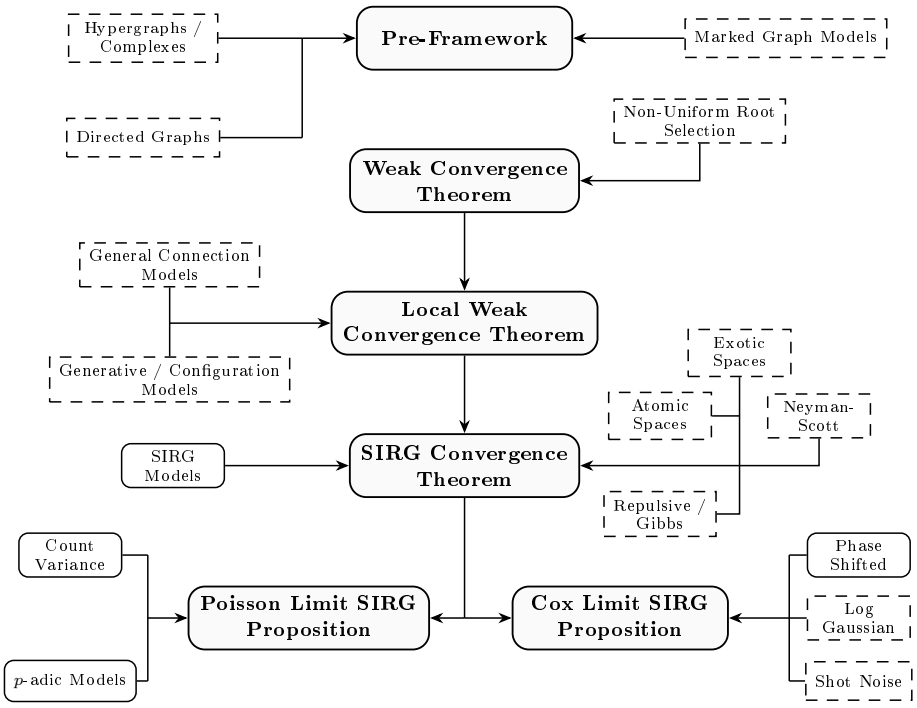

Figure 2: Overview of the main convergence results and model families feeding into them. Dashed borders indicate that the topic was not directly explored in this paper.

At the lowest level, we have the limits of the framework itself. Although our adjacency processes need not be symmetric, the graph construction map collapses this asymmetry into an ultimately undirected resulting graph. This means that if one wants to analyze directed graphs, multigraphs, or hypergraphs, they would not be able to directly utilize our framework as is. However, it is likely that proving analogous versions of our results for these models is viable using similar measure-theoretic arguments and intuition. Marked graphs also fit here. While we can handle some types of marking by treating vertex marks as part of the label space, handling edge marking is more difficult and not transparent via our theory.

The first truly usable level is thus our weak convergence theorem, Theorem 4.1, and the surrounding content in Section 4. However, by far the most common type of convergence analysis on random graphs is local weak convergence, and therefore most models can plug directly into Section 5. The only models that truly fit at this level are therefore models where weak convergence is taken against a different rooting mechanism, or possibly models where rooted graphs themselves are the fundamental objects.

As soon as local weak convergence becomes the desired convergence mechanism,

one is better off instead using our local weak convergence theorem, Theorem 5.2, and the surrounding machinery in Section 5. This is where one would go if they had a general connection model, possibly with dependent edge relations, that did not fit within the SIRG umbrella. More exotic random graph models that are constructed via iterative generative rules, such as preferential-attachment models [4], or via degree-sequence configuration mechanisms whose local weak limits are unimodular/size-biased Galton-Watson type trees [17, Sec. 4.2], should also use this level as a starting interface.

It should also be noted that it is technically possible to explore models in this area directly using the vertex-normalized Palm distributions, although this is far less standard than classical Palm theory, and therefore relying on vertex concentration is advised for most models.

If one has a SIRG model, they can leverage the full results of Section 6, including the local weak convergence theorem for SIRGs, Theorem 6.4. This is where one would go if they want to understand the local weak behavior of a class of models that is not covered by our selection of examples. This would most likely come in one of two ways. Firstly, one might want to explore more complicated vertex processes on $\mathbb{R}^d$ or $\mathbb{Q}_p^d$, such as Neyman-Scott style clustering processes [31], determinantal/repulsive style processes [27], or general Gibbs processes [28]. Secondly, one might want to explore vertices in more exotic spaces, although real difficulty appears here as finding β-admissible sequences for such spaces is quite difficult. The main class of approachable spaces would be atomic spaces [20], on which one might want to explore Bernoulli site processes [14] or other related discrete processes. In general though, the most viable path to extension would be more intricate vertex processes such as those aforementioned, as including these likely amounts to bounding the higher moment distributions used in the conditions of Theorem 6.4.

It should also be noted that even for the models that cannot directly utilize the results of Section 6, it is likely that the results from this section can still provide insight on proof strategies and model analysis.

If one is curious about Poisson or Cox vertex processes, they can directly leverage the results in Section 7, although some work is still needed to verify directing measure convergence for more complicated directing measures, such as log Gaussian Cox processes (LGCPs) and shot-noise Cox processes (SNCPs) [15] within the scope of Proposition 7.25.

With the results reviewed, we move on to a more in-depth discussion about limitations of the theory and viable future research directions.

8.2 Limitations and Potential Improvements

The potential extensions of this framework are mostly not structural, but are rather improvements to the generality of conditions and verification of more model classes. However, before proceeding to these, there is one structural issue

that is worth noting. This issue is that in general, constructing this theory via directed graphs, with undirected graphs as a subclass (by identifying 2-cycles with undirected edges), is probably the cleaner way to represent this theory.

This is due to the fact that adjacency processes are asymmetric in nature, while the graph construction map is what collapses this asymmetry into the undirected Benjamini-Schramm topology. In a sense, this is analogous to analyzing a complex function by collapsing all of its outputs into their real part. This does not discredit the present framework, but rather highlights its limitations. For example, while we allow asymmetric adjacency measures, any model one would care to use effectively needs to be symmetric, in order to preserve true edge densities in the resulting graph and to allow the study of distributions of graph observables as in [18]. This was exactly the motivation for making SIRGs symmetric in Section 6.

Thus, extensions to the present framework would require strictly symmetric modelling in order to retain the aforementioned properties. Alternatively, one could refine the present theory to construct graphs in the directed equivalent of the Benjamini-Schramm space (with balls taken in the underlying undirected graph), and treat undirected SIRGs as a special symmetric variant within a broader class of SIRGs, from which Section 6 style machinery could be developed. However, this directed reformulation is not a modular add-on to the present undirected map. Once the graph-construction map and neighborhood objects are changed, one must restate and re-prove essentially all core lemmas and propositions in the directed setting. Even if most arguments remain structurally analogous and the new proofs are mostly technical refinements, this still creates substantial duplicated verification effort. For that reason, a directed graph convention is likely the better foundational baseline in the long run. Otherwise, fixing the framework at symmetric restrictions would risk forcing a near-complete reconstruction later when generalizing to the directed map.

Regardless of the choice of future refinement, the present framework still stands strong as a central pipeline for determining local weak convergence and an isolation of necessary criteria, and would undoubtedly be useful as a baseline on which to develop one of the above refinements.

With the main structural issue out of the way, we move to covering the main areas where there is room for improvement and generalization. The first topic we will cover is that of radial tightness and fibre-finiteness. Both of these conditions are very difficult to show in general. We were only able to find concrete criteria for these conditions by considerably and repeatedly restricting model classes. The smaller the model class, the more checkable this condition is. However, it is likely that the radial tightness condition given in Lemma 6.7 is itself not maximally general, and similarly for the fibre-finiteness condition in Lemma 6.6. A very useful future research direction would be to refine these conditions and to make them simpler to check and more general, as this would aid in the ultimate goal of providing a modular framework for easily identifying local weak convergence.

In addition to this, weak convergence of graphs is also somewhat limited in its capacity. A stronger type of convergence is known as local convergence in probability, see [18, Def. 1.4]. Fortunately, determining such a convergence is a modular extension and would not require a rewrite of the framework. It would instead be an additional condition to check at each level of our theory, as local convergence in probability implies local weak convergence, and thus there is no way to conclude local convergence in probability without at least the machinery of this framework. While it is modular at the level of framework architecture, the improvement would still be technically non-trivial, typically requiring additional concentration/second-moment control. A future extension could determine the specific conditions under which local weak convergence can be upgraded to this stronger version at each of our levels, which would make the results considerably more applicable.

Another issue is convergence of marking kernels. To an extent this ties into the limitations of the current radial tightness criterion, Lemma 6.7, as it is the only condition preventing us from using a Proposition 6.5 style marking kernel convergence in the general SIRG theorem. An extension here is likely possible, and would align our results better with the existing theory in [18].

Of course, there are also several strong research directions that are not covered in depth here. In general, any extensions seen in Figure 2 are viable to explore, along with extensions in the same spirit. Results on computing the distributions of observables such as finite cycles and triangles are also essential, but such an extension would likely be fruitful only on symmetric models or after the aforementioned structural issues are patched.

Taken together, these limitations do not overturn the central outcome of the paper. The framework provides a coherent and reusable route from abstract convergence theory to model-level verification, and the examples indicate that it can be applied in practice beyond purely formal statements. While some components can still be refined, the core machinery is stable, useful, and a strong basis for further work.

A Proofs

A.1 Lemma 3.1: Measurability of Vertex Map

Proof. Let

$$\Delta^2(S) := \{(x, x) : x \in S\} \subseteq S^2$$

denote the diagonal and, for $B \in \mathcal{B}(S)$, write

$$\Delta^2(B) := \{(x, x) : x \in B\} \subseteq S^2.$$

The diagonal map

$$\iota : S \to S^2, \qquad \iota(x) = (x, x),$$

is a homeomorphism [29, Def. 4.2.1] from S onto the closed subspace $\Delta^2(S) \subseteq S^2$. Hence, if $B \subseteq S$ is Borel, then $\Delta^2(B) = \iota(B)$ is Borel in S^2. Moreover, if B is relatively compact, then $\Delta^2(B)$ is relatively compact in S^2, since

$$\overline{\Delta^2(B)} \subseteq \Delta^2(\overline{B}),$$

and $\Delta^2(\overline{B})$ is compact as the continuous image of the compact set $\overline{B}$. Now let $B \subseteq S$ be relatively compact and Borel. By definition of the vertex map,

$$V(\tau) = \sum_{(x,x)\in\tau} \delta_x.$$

Therefore, for every $\tau \in \mathcal{A}(S^2)$,

$$V(\tau)(B) = \sum_{(x,x)\in\tau} \delta_x(B) = \tau(\Delta^2(B)).$$

Thus the evaluation map

$$\tau \longmapsto V(\tau)(B)$$

is the restriction to $\mathcal{A}(S^2)$ of the evaluation map

$$\mu \longmapsto \mu(\Delta^2(B))$$

on $\mathcal{N}(S^2)$. Since $\Delta^2(B)$ is relatively compact and Borel, this evaluation map is measurable for the vague Borel σ-algebra on $\mathcal{N}(S^2)$. Hence,

$$\tau \longmapsto V(\tau)(B)$$

is measurable on $\mathcal{A}(S^2)$. The Borel σ-algebra on $\mathcal{N}(S)$ (with the vague topology) is generated [30, p. 28] by the evaluation maps $\nu \mapsto \nu(B)$, where $B \subseteq S$ ranges over relatively compact Borel sets. Since each such coordinate map is measurable after composition with V, the map

$$V : \mathcal{A}(S^2) \longrightarrow \mathcal{N}(S)$$

is measurable. ◀

A.2 Lemma 3.2: Adjacency Measure Spaces are Polish

Proof. Throughout, $\mathcal{N}(S^2)$ denotes the space of boundedly finite simple counting measures on S^2, equipped with the vague topology. Since S^2 is LC Polish, $\mathcal{N}(S^2)$ is Polish [30, p. 28].

We use the following standard facts about vague convergence of simple counting measures. If $\nu_n \to \nu$ in $\mathcal{N}(E)$, where E is LC Polish, then:

1. if $z \in \nu$, then along a subsequence there exist atoms $z_n \in \nu_n$ such that $z_n \to z$;
2. if $z_n \in \nu_n$ and $z_n \to z$, then $z \in \nu$.

These both follow directly from standard atom matching behavior [2, Prop. 2.8]. We first prove that $\mathcal{A}(S^2)$ is Polish. It suffices to show that $\mathcal{A}(S^2)$ is closed in $\mathcal{N}(S^2)$. Let $\tau_n \in \mathcal{A}(S^2)$ and suppose that

$$\tau_n \to \tau$$

vaguely in $\mathcal{N}(S^2)$. We show that τ satisfies the adjacency condition. Let $(x, y) \in \tau$. By localization of atoms, after passing to a subsequence there exist atoms $(x_n, y_n) \in \tau_n$ such that

$$(x_n, y_n) \to (x, y).$$

Since each τ_n is an adjacency measure,

$$(x_n, x_n) \in \tau_n, \qquad (y_n, y_n) \in \tau_n.$$

By continuity of the diagonal map $z \mapsto (z, z)$,

$$(x_n, x_n) \to (x, x), \qquad (y_n, y_n) \to (y, y).$$

The closedness of atoms under vague convergence therefore gives

$$(x, x) \in \tau, \qquad (y, y) \in \tau.$$

Hence $\tau \in \mathcal{A}(S^2)$. Thus $\mathcal{A}(S^2)$ is closed in $\mathcal{N}(S^2)$, and therefore Polish [35, Prop. 1.4]. Next consider the rooted space. By definition,

$$\mathcal{A}^\circ(S^2) = \{(\tau, x) \in \mathcal{A}(S^2) \times S : (x, x) \in \tau\}.$$

Since $\mathcal{A}(S^2)$ and S are Polish, the product $\mathcal{A}(S^2) \times S$ is Polish. We show that $\mathcal{A}^\circ(S^2)$ is closed in this product.

Suppose

$$(\tau_n, x_n) \to (\tau, x)$$

in $\mathcal{A}(S^2) \times S$, with $(x_n, x_n) \in \tau_n$ for every n. Then

$$(x_n, x_n) \to (x, x).$$

By closedness of atoms under vague convergence,

$$(x, x) \in \tau.$$

Therefore $(\tau, x) \in \mathcal{A}^\circ(S^2)$. Hence $\mathcal{A}^\circ(S^2)$ is closed in $\mathcal{A}(S^2) \times S$, and is Polish. Now define the coordinate-swap map

$$\rho : S^2 \to S^2, \qquad \rho(x, y) = (y, x).$$

Since ρ is a homeomorphism, the pushforward map

$$\tau \mapsto \rho_\# \tau$$

is continuous under the vague topology. Indeed, for every $f \in \mathcal{C}_c(S^2)$,

$$\int f \, d(\rho_\# \tau) = \int f \circ \rho \, d\tau,$$

and $f \circ \rho \in \mathcal{C}_c(S^2)$. The symmetric subspace can be written as

$$\mathcal{A}_{\mathrm{sym}}(S^2) = \{\tau \in \mathcal{A}(S^2) : \rho_\# \tau = \tau\}.$$

To see that this set is closed, let $\tau_n \in \mathcal{A}_{\mathrm{sym}}(S^2)$ with $\tau_n \to \tau$ in $\mathcal{A}(S^2)$. By continuity of the pushforward map $\rho_\#$, we have $\rho_\# \tau_n \to \rho_\# \tau$. But $\rho_\# \tau_n = \tau_n$ for every n, and $\tau_n \to \tau$. Since $\mathcal{N}(S^2)$ is Hausdorff, limits are unique, so

$$\rho_\# \tau = \tau.$$

Thus, $\tau \in \mathcal{A}_{\mathrm{sym}}(S^2)$, proving that $\mathcal{A}_{\mathrm{sym}}(S^2)$ is closed in $\mathcal{A}(S^2)$. Hence, it is Polish. Finally,

$$\mathcal{A}^\circ_{\mathrm{sym}}(S^2) = \mathcal{A}^\circ(S^2) \cap \big(\mathcal{A}_{\mathrm{sym}}(S^2) \times S\big).$$

Since $\mathcal{A}_{\mathrm{sym}}(S^2)$ is closed in $\mathcal{A}(S^2)$, the set $\mathcal{A}_{\mathrm{sym}}(S^2) \times S$ is closed in $\mathcal{A}(S^2) \times S$. Therefore, $\mathcal{A}^\circ_{\mathrm{sym}}(S^2)$ is closed in $\mathcal{A}^\circ(S^2)$, and is Polish. It remains to prove measurability of the symmetrization maps. Let $\mathcal{L}_{\mathrm{ct}}(S^2)$ denote the space of boundedly finite counting measures on S^2, not necessarily simple. Define

$$T : \mathcal{A}(S^2) \to \mathcal{L}_{\mathrm{ct}}(S^2), \qquad T(\tau) := \tau + \rho_\# \tau.$$

The map T is continuous, since $\tau \mapsto \rho_\# \tau$ is continuous and addition of boundedly finite measures is continuous under the vague topology. Define the support-reduction map

$$R : \mathcal{L}_{\mathrm{ct}}(S^2) \to \mathcal{N}(S^2), \qquad R(\mu) := \sum_{z : \mu(\{z\}) > 0} \delta_z.$$

We show that R is measurable. Let $U \subseteq S^2$ be relatively compact and open. Fix a countable basis $\mathcal{B}$ for the topology of S^2. For every $k \in \mathbb{N}$, the event

$$\{R(\mu)(U) \geq k\}$$

says that the support of μ contains at least k distinct points in U. Equivalently, there exist pairwise disjoint basis elements $B_1, \dots, B_k \in \mathcal{B}$ with $B_i \subseteq U$ such that

$$\mu(B_i) \geq 1 \qquad \text{for all } i = 1, \dots, k.$$

Hence,

$$\{R(\mu)(U) \geq k\} = \bigcup_{\substack{B_1,\dots,B_k \in \mathcal{B} \\ B_i \subseteq U,\ B_i \cap B_j = \varnothing}} \bigcap_{i=1}^{k} \{\mu(B_i) \geq 1\}.$$

The right-hand side is measurable, since the evaluation maps $\mu \mapsto \mu(B_i)$ are measurable. Therefore, $\mu \mapsto R(\mu)(U)$ is measurable for every relatively compact open U. Since the Borel σ-algebra on $\mathcal{N}(S^2)$ is generated by such counting maps, R is measurable. Now observe that

$$\mathrm{Sym}(\tau) = R\big(\tau + \rho_{\#}\tau\big) = (R \circ T)(\tau).$$

Since T is continuous and R is measurable, Sym is measurable as a map

$$\mathrm{Sym} : \mathcal{A}(S^2) \to \mathcal{N}(S^2).$$

We check that its image lies in $\mathcal{A}_{\mathrm{sym}}(S^2)$. Let $\tau \in \mathcal{A}(S^2)$. By construction,

$$(x, y) \in \mathrm{Sym}(\tau) \quad \Longleftrightarrow \quad (x, y) \in \tau \text{ or } (y, x) \in \tau.$$

Therefore,

$$(x, y) \in \mathrm{Sym}(\tau) \quad \Longleftrightarrow \quad (y, x) \in \mathrm{Sym}(\tau),$$

so $\mathrm{Sym}(\tau)$ is symmetric. It remains only to verify that $\mathrm{Sym}(\tau)$ is still an adjacency measure. It is simple by definition of R. It is boundedly finite because, for every bounded Borel set $B \subseteq S^2$,

$$\mathrm{Sym}(\tau)(B) \leq \tau(B) + \rho_{\#}\tau(B) = \tau(B) + \tau(\rho^{-1}(B)),$$

and both terms are finite. Finally, if $(x, y) \in \mathrm{Sym}(\tau)$, then either $(x, y) \in \tau$ or $(y, x) \in \tau$. Since τ is an adjacency measure, in both cases

$$(x, x) \in \tau, \qquad (y, y) \in \tau.$$

Hence

$$(x, x), (y, y) \in \mathrm{Sym}(\tau).$$

Thus

$$\mathrm{Sym}(\tau) \in \mathcal{A}_{\mathrm{sym}}(S^2).$$

Since $\mathcal{A}_{\mathrm{sym}}(S^2)$ carries the subspace Borel σ-algebra inherited from $\mathcal{A}(S^2)$, the map

$$\mathrm{Sym} : \mathcal{A}(S^2) \to \mathcal{A}_{\mathrm{sym}}(S^2)$$

is measurable.

Finally, define

$$\mathrm{Sym}^\circ : \mathcal{A}^\circ(S^2) \to \mathcal{A}^\circ_{\mathrm{sym}}(S^2), \qquad \mathrm{Sym}^\circ(\tau, x) := (\mathrm{Sym}(\tau), x).$$

This map is measurable as the product of the measurable map Sym and the identity map on S. Moreover, if $(\tau, x) \in \mathcal{A}^\circ(S^2)$, then $(x, x) \in \tau$, and therefore

$$(x, x) \in \mathrm{Sym}(\tau).$$

Thus,

$$\mathrm{Sym}^\circ(\tau, x) \in \mathcal{A}^\circ_{\mathrm{sym}}(S^2).$$

Since $\mathcal{A}^\circ_{\mathrm{sym}}(S^2)$ carries the inherited subspace Borel σ-algebra, Sym° is measurable. ◀

A.3 Lemma 3.5: Measurability of Graph Construction

Proof. Let $(\tau, x) \in \mathcal{E}^\circ(S^2)$. By definition, τ is fiber-finite and $x \in V_\tau$. Hence, $G_S(\tau)$ is locally finite by Lemma 3.3. Therefore,the connected component $\mathcal{C}_{G_S(\tau)}(x)$ is a connected locally finite graph rooted at x, and so

$$G(\tau, x) = [(\mathcal{C}_{G_S(\tau)}(x), x)] \in \mathcal{G}_*.$$

Thus, the map is well-defined.

It remains to prove measurability. We use the following standard measurable facts. First, the atom relation

$$\{(\mu, z) \in \mathcal{N}(X) \times X : \mu(\{z\}) = 1\}$$

is Borel for every LC Polish space X. Indeed, the map $(\mu, z) \mapsto \mu(\{z\})$ is measurable since

$$\mu(\{z\}) = \int_X 1_{\Delta^2(X)}(z, y)\, \mu(dy),$$

where $\Delta^2(X) = \{(z, z) \in X^2 : z \in X\}$ is Borel because X is Polish. Hence $\{(\mu, z) : \mu(\{z\}) = 1\}$ is Borel. Applied with $X = S^2$, this implies that

$$\{(\tau, u, v) \in \mathcal{E}^\circ(S^2) \times S \times S : (u, v) \in \tau\}$$

is Borel. Secondly, every Polish space admits a Borel strict total order. Concretely, every Polish space can be Borel embedded into $[0, 1]$, by [5, Thm. 6.7.4]. Pulling back the usual order on $\mathbb{R}$ gives a strict total order $\prec$ on Y: for $y_1, y_2 \in Y$,

$$y_1 \prec y_2 \quad \iff \quad \phi(y_1) < \phi(y_2),$$

where $\phi : Y \to \mathbb{R}$ is a Borel injective map. The order is called Borel because

$$\{(y_1, y_2) \in Y^2 : y_1 \prec y_2\}$$

is a Borel subset of Y^2. After fixing such an order, every non-empty finite subset of Y has a unique $\prec$-maximal element. Moreover, if M is a measurable finite

set-valued map into Y, then selecting this maximal element is measurable on $\{M \neq \varnothing\}$. Thus the order provides a measurable tie-breaking rule for finite sets. Fix such a Borel strict total order $\prec$ on S. Adjoin an isolated cemetery point $\partial \notin S$, and write

$$\bar{S} := S \cup \{\partial\}.$$

We construct a measurable enumeration

$$o_0, o_1, o_2, \cdots : \mathcal{E}^\circ(S^2) \to \bar{S}$$

of the connected component of the root. The enumeration is only used as a measurable coding of the rooted component; it is not part of the graph structure. Set

$$o_0(\tau, x) := x.$$

Suppose $o_0, \ldots, o_n$ have already been constructed measurably. Define the explored set

$$F_n(\tau, x) := \{o_0(\tau, x), \ldots, o_n(\tau, x)\} \setminus \{\partial\}.$$

Define the frontier

$$N_n(\tau, x) := \left\{ z \in S \setminus F_n(\tau, x) \;\middle|\; \begin{array}{l} \exists\, k \leq n \text{ such that } o_k(\tau, x) \neq \partial, \\ \big(o_k(\tau, x), z\big) \in \tau \text{ or } \big(z, o_k(\tau, x)\big) \in \tau \end{array} \right\}.$$

Since $F_n(\tau, x)$ is finite and τ is fiber-finite, $N_n(\tau, x)$ is finite. The membership relation

$$\{((\tau, x), z) : z \in N_n(\tau, x)\}$$

is Borel. Indeed, it is obtained from the already measurable maps $o_0, \ldots, o_n$, the Borel atom relation

$$\{(\tau, u, v) : (u, v) \in \tau\},$$

and finitely many Boolean operations. Hence, N_n is a measurable finite set-valued map. Define

$$o_{n+1}(\tau, x) := \begin{cases} \max_{\prec} N_n(\tau, x), & N_n(\tau, x) \neq \varnothing, \\ \partial, & N_n(\tau, x) = \varnothing. \end{cases}$$

By the measurable selection fact, o_{n+1} is measurable. Inductively, all maps o_n are measurable. This enumeration lists exactly the vertices in the connected component of x. Indeed, every selected point is adjacent to an already selected point, so each $o_n \neq \partial$ lies in $\mathcal{C}_{G_S(\tau)}(x)$. Conversely, if a vertex z in this component was never selected, take a shortest path from x to z. Along this path, there is a first vertex not selected; its predecessor is selected, so it eventually belongs to a non-empty frontier N_n, contradicting the fact that the construction only stops when the frontier is empty. For $u, v \in \bar{S}$, write

$$u \sim_\tau v \quad \Longleftrightarrow \quad u, v \in S, \quad u \neq v, \quad (u, v) \in \tau \text{ or } (v, u) \in \tau.$$

For $i, j \in \mathbb{N}$, define the adjacency coordinates

$$a_{ij}(\tau, x) := \mathbf{1}_{\{o_i(\tau,x) \sim_\tau o_j(\tau,x)\}}.$$

These maps are measurable because the maps o_i, o_j are measurable and the relation

$$\{(\tau, u, v) : (u, v) \in \tau\}$$

is Borel. The sequence $(a_{ij})_{i,j\in\mathbb{N}}$, together with the distinguished index 0, determines the rooted isomorphism class

$$[\mathcal{C}_{G_S(\tau)}(x), x].$$

We now use this coding to verify measurability with respect to the canonical σ-algebra on $\mathcal{G}_*$. The Borel σ-algebra on $\mathcal{G}_*$ is generated by the rooted ball cylinder sets

$$C_{r,\tilde{G}} := \{[G, o] \in \mathcal{G}_* : B_r(G, o) \cong \tilde{G}\},$$

where $r \in \mathbb{N}$ and $\tilde{G}$ ranges over finite rooted graphs [18, Def. 1.2]. Therefore, it suffices to show that

$$G^{-1}(C_{r,\tilde{G}})$$

is measurable for every such r and $\tilde{G}$. For each $n \in \mathbb{N}$, define the event that the listed vertex o_n lies within graph distance at most r from the root by

$$D_{n,r} := \left\{ (\tau, x) \;\middle|\; \begin{array}{l} \exists\, m \le r,\ \exists\, i_0, \ldots, i_m \in \mathbb{N} \text{ such that} \\ i_0 = 0, \qquad i_m = n, \\ a_{i_\ell i_{\ell+1}}(\tau, x) = 1 \quad \text{for all } \ell = 0, \ldots, m-1 \end{array} \right\}.$$

This is measurable, since it is a countable union of finite intersections of measurable coordinate events. It represents the event

$$d_{G_S(\tau)}(x, o_n(\tau, x)) \le r.$$

Since $G_S(\tau)$ is locally finite, the radius-r ball around x contains only finitely many vertices. Hence, for a fixed finite set $I \subset \mathbb{N}$, the event

$$I = \{n \in \mathbb{N} : D_{n,r} \text{ holds}\}$$

is measurable, because it equals

$$\left(\bigcap_{n \in I} D_{n,r} \right) \cap \left(\bigcap_{n \notin I} D_{n,r}^c \right),$$

a countable intersection of measurable sets. Now fix r and a finite rooted graph $\tilde{G}$. The event $G^{-1}(C_{r,\tilde{G}})$ can be written as the countable union, over all finite $I \subset \mathbb{N}$, of the event that:

1. $I = \{n : D_{n,r} \text{ holds}\}$, and

2. the finite graph on vertex set I, rooted at 0, with adjacency matrix $(a_{ij})_{i,j\in I}$, is rooted-isomorphic to $\tilde{G}$.

The first condition is measurable by the previous paragraph. The second condition depends only on finitely many coordinates a_{ij}, $i, j \in I$, and is therefore measurable. Since there are only countably many finite subsets $I \subset \mathbb{N}$, the union is measurable. Thus, $G^{-1}(C_{r,\tilde{G}})$ is measurable for every rooted-ball cylinder $C_{r,\tilde{G}}$. Since these cylinders generate the canonical σ-algebra on $\mathcal{G}_*$, the map

$$G : \mathcal{E}^\circ(S^2) \to \mathcal{G}_*$$

is measurable. ◀

A.4 Lemma 3.6: Measurability of Other Graph Maps

Proof. We take measurability of the rooted-ball map

$$B_r : \mathcal{G}_* \to \mathcal{G}_*$$

as established [17, Def. A.3]. It remains to prove well-definedness and measurability of

$$V_r : \mathcal{E}^\circ(S^2) \to \mathcal{N}(S).$$

Let $(\tau, x) \in \mathcal{E}^\circ(S^2)$. Since τ is regular, it is fiber-finite. Hence, by Lemma 3.3, the constructed labelled graph $G_S(\tau)$ is locally finite. Therefore, the radius-r graph ball around x contains only finitely many vertices. Define

$$W_r(\tau, x) := \{y \in V_\tau : d_{G_S(\tau)}(x, y) \le r\}.$$

Then $W_r(\tau, x)$ is finite, and therefore

$$V_r(\tau, x) := \sum_{y \in W_r(\tau,x)} \delta_y$$

is a boundedly finite simple counting measure on S. Hence $V_r(\tau, x) \in \mathcal{N}(S)$, so V_r is well-defined. We now prove measurability. By the proof of Lemma 3.5, there exist measurable maps

$$o_0, o_1, o_2, \cdots : \mathcal{E}^\circ(S^2) \to \bar{S}$$

enumerating the connected component of the root, together with measurable adjacency coordinates

$$a_{ij} : \mathcal{E}^\circ(S^2) \to \{0, 1\}.$$

For $n \in \mathbb{N}$, define

$$D_{n,r} := \left\{ (\tau, x) \;\middle|\; \begin{array}{l} \exists\, m \le r,\ \exists\, i_0, \ldots, i_m \in \mathbb{N} \text{ such that} \\ i_0 = 0, \qquad i_m = n, \\ a_{i_\ell i_{\ell+1}}(\tau, x) = 1 \quad \text{for all } \ell = 0, \ldots, m-1 \end{array} \right\}.$$

Each $D_{n,r}$ is measurable, because it is a countable union of finite intersections of measurable coordinate events. Moreover, $D_{n,r}$ is exactly the event that the listed vertex $o_n(\tau, x)$ lies at graph distance at most r from the root. Thus, for every Borel set $A \subseteq S$,

$$V_r(\tau, x)(A) = \sum_{n \geq 0} \mathbf{1}_{D_{n,r}}(\tau, x) \mathbf{1}_{\{o_n(\tau,x) \in A\}}.$$

Since $G_S(\tau)$ is locally finite, only finitely many vertices lie within graph distance r of x. Hence the above sum is pointwise finite. Each summand is measurable, so

$$(\tau, x) \mapsto V_r(\tau, x)(A)$$

is measurable for every Borel set $A \subseteq S$. Since the Borel σ-algebra on $\mathcal{N}(S)$ is generated by the evaluation maps $\nu \mapsto \nu(A)$, it follows that

$$V_r : \mathcal{E}^\circ(S^2) \to \mathcal{N}(S)$$

is measurable. Therefore all maps from Definition 12 are well-defined and measurable on their respective Borel σ-algebras. ◀

A.5 Lemma 3.7: Uniform Rooting Kernel

Proof. Let

$$\mathcal{L} := \{\tau \in \mathcal{E}(S^2) : \tau(S^2) < \infty\}$$

be the subspace of finite adjacency measures. We must show that

$$\mathbb{U} : \mathcal{L} \times \mathcal{B}(S) \to [0, 1]$$

is a Markov kernel and that it selects roots from the vertex set. First fix $\tau \in \mathcal{L}$. Since τ is finite and non-empty, its vertex set $V(\tau)$ is finite and non-empty. Therefore,

$$1 \leq |V(\tau)| < \infty.$$

Hence, for every $A \in \mathcal{B}(S)$,

$$\mathbb{U}(\tau, A) = \frac{|V(\tau) \cap A|}{|V(\tau)|}$$

is well-defined. Moreover, the map

$$A \mapsto \mathbb{U}(\tau, A)$$

is a probability measure on $(S, \mathcal{B}(S))$: it is the normalized counting measure on the finite set $V(\tau)$. Indeed,

$$\mathbb{U}(\tau, S) = \frac{|V(\tau)|}{|V(\tau)|} = 1,$$

and countable additivity follows from countable additivity of the counting measure, or directly from finiteness of $V(\tau)$. It remains to show that, for every $A \in \mathcal{B}(S)$, the map

$$\tau \mapsto \mathbb{U}(\tau, A)$$

is measurable on $\mathcal{L}$. Recall that the vertex map

$$V : \mathcal{E}(S^2) \to \mathcal{N}(S)$$

is measurable by Lemma 3.1. Therefore, the maps

$$\tau \mapsto V(\tau)(A) \qquad \text{and} \qquad \tau \mapsto V(\tau)(S)$$

are measurable on $\mathcal{L}$. Since

$$V(\tau)(A) = |V(\tau) \cap A|, \qquad V(\tau)(S) = |V(\tau)|,$$

we have

$$\mathbb{U}(\tau, A) = \frac{V(\tau)(A)}{V(\tau)(S)}.$$

The denominator is strictly positive on $\mathcal{L}$, because every $\tau \in \mathcal{L}$ is non-empty and regular. Hence, $\tau \mapsto \mathbb{U}(\tau, A)$ is measurable. Finally, $\mathbb{U}$ is supported on the vertex set. Indeed,

$$\mathbb{U}(\tau, S \setminus V(\tau)) = \frac{|V(\tau) \cap (S \setminus V(\tau))|}{|V(\tau)|} = 0.$$

Thus, $\mathbb{U}$ is a Markov kernel from $\mathcal{L}$ to S and satisfies the defining support condition of a rooting kernel. Therefore, $\mathbb{U}$ is a rooting kernel. ◀

A.6 Lemma 7.14: Radial Growth Domination

Proof. Write for clarity

$$m_n := \lambda_{(S,H)}(S_n), \qquad \lambda_n := \lambda_{(S,H)}|_{S_n}, \qquad Y_n := \frac{N_n}{m_n},$$

We verify condition (5) of Theorem 6.4 first. By Equation (31),

$$\Gamma_n^x \overset{d}{=} \delta_x + \Gamma_n^{!x}.$$

By Lemma 7.8,

$$\Lambda_{\Gamma_n^{!x}} = a_{1,n}\lambda_n, \qquad \alpha^{(2)}_{\Gamma_n^{!x}} = a_{2,n}\lambda_n^{\otimes 2},$$

where

$$a_{1,n} := \frac{\mathbb{E}[(N_n)_2]}{\mathbb{E}[N_n]\, m_n}, \qquad a_{2,n} := \frac{\mathbb{E}[(N_n)_3]}{\mathbb{E}[N_n]\, m_n^2}.$$

By Assumptions 7.7 and 7.10, and as shown in the proof of Lemma 7.12, we have

$$\mathbb{E}[Y_n] \to c, \qquad \mathbb{E}[Y_n^2] \to c^2, \qquad \sup_n \mathbb{E}[Y_n^3] < \infty.$$

Therefore,

$$a_{1,n} = \frac{\mathbb{E}[Y_n^2] - \mathbb{E}[Y_n]/m_n}{\mathbb{E}[Y_n]}, \qquad a_{2,n} = \frac{\mathbb{E}[Y_n^3] - 3\mathbb{E}[Y_n^2]/m_n + 2\mathbb{E}[Y_n]/m_n^2}{\mathbb{E}[Y_n]},$$

so $(a_{1,n})_{n\in\mathbb{N}}$ and $(a_{2,n})_{n\in\mathbb{N}}$ are eventually bounded. Since $\Gamma_n^x = \delta_x + \Gamma_n^{!x}$, its second factorial moment measure decomposes as

$$\alpha^{(2)}_{\Gamma_n^x} = \alpha^{(2)}_{\Gamma_n^{!x}} + \delta_x \otimes \Lambda_{\Gamma_n^{!x}} + \Lambda_{\Gamma_n^{!x}} \otimes \delta_x.$$

Fix a compact set $L \subset S$, let $O_R := S \setminus B_S(x, R)$, and define

$$I_{n,R}(L) := \iint_{L\times O_R} \overline{\kappa}_n(d_{S_n}(u,v))\, \alpha^{(2)}_{\Gamma_n^x}(du, dv).$$

Since $x \notin O_R$, the term $\Lambda_{\Gamma_n^{!x}} \otimes \delta_x$ does not contribute, so

$$\begin{aligned} I_{n,R}(L) = a_{2,n} \iint_{(L\cap S_n)\times(O_R\cap S_n)} \overline{\kappa}_n(d_{S_n}(u,v))\, \lambda_n(du)\lambda_n(dv) \\ + \mathbf{1}_{\{x\in L\}}\, a_{1,n} \int_{O_R\cap S_n} \overline{\kappa}_n(d_{S_n}(x,v))\, \lambda_n(dv). \end{aligned} \tag{43}$$

Set

$$J^{(2)}_{n,R}(L) := \iint_{(L\cap S_n)\times(O_R\cap S_n)} \overline{\kappa}_n(d_{S_n}(u,v))\, \lambda_n(du)\lambda_n(dv), \tag{44}$$

$$J^{(1)}_{n,R} := \int_{O_R\cap S_n} \overline{\kappa}_n(d_{S_n}(x,v))\, \lambda_n(dv). \tag{45}$$

Then Equation (43) becomes

$$I_{n,R}(L) = a_{2,n} J^{(2)}_{n,R}(L) + \mathbf{1}_{\{x\in L\}}\, a_{1,n} J^{(1)}_{n,R}.$$

By Assumption 7.13, there exist $N_0 \in \mathbb{N}$ and $t_0 > 0$ such that

$$\overline{\kappa}_n(t) \le \phi(t) \qquad \text{for all } n \ge N_0 \text{ and } t \ge t_0.$$

We then use the same β-admissibility compact-containment argument as in the proof of Lemma 6.7, since β-admissibility is part of the context. By conditions (2)-(3) of Definition 23, there exists some $N_1 > 0$ and $1 > \delta > 0$ satisfying for $n \ge N_1$

$$d_{S_n}(x,v) > d_S(x,v) - \delta.$$

Hence, if $n \ge \max\{N_0, N_1\}$ and $R \ge t_0 + \delta$, then

$$\overline{\kappa}_n(d_{S_n}(x,v)) \le \phi(d_{S_n}(x,v)) \le \phi(d_S(x,v) - \delta).$$

Hence,

$$J_{n,R}^{(1)} \leq \int_{O_R} \phi(d_S(x,v) - \delta)\, \lambda_{(S,H)}(dv) = \int_R^\infty \phi(t-\delta)\, dV_x(t). \tag{46}$$

For the two-point term, set

$$r_L := \sup_{u \in L} d_S(x,u) < \infty.$$

If $u \in L$ and $v \in O_R$, then

$$d_S(u,v) \geq d_S(x,v) - d_S(x,u) > R - r_L.$$

Again by the same β-admissibility argument, there exists some $N_1 > 0$ and $1 > \delta > 0$ satisfying for $n \geq N_1$

$$d_{S_n}(u,v) > d_S(u,v) - \delta$$

for every $u \in L \cap S_n$, $v \in O_R \cap S_n$, and. Therefore, if $R \geq t_0 + r_L + \delta$, then for every such u, v and every $n \geq \max\{N_0, N_1\}$,

$$\overline{\kappa}_n(d_{S_n}(u,v)) \leq \phi(d_{S_n}(u,v)) \leq \phi(d_S(u,v) - \delta).$$

Thus,

$$\begin{aligned}
\int_{O_R \cap S_n} \overline{\kappa}_n(d_{S_n}(u,v))\, \lambda_{(S,H)}(dv) &\leq \int_{O_R} \phi(d_S(u,v) - \delta)\, \lambda_{(S,H)}(dv) \\
&\leq \int_{R-r_L}^\infty \phi(t-\delta)\, dV_u(t) \\
&= \int_{R-r_L}^\infty \phi(t-\delta)\, dV_x(t)
\end{aligned}$$

where the last step uses transitive isometric invariance. Integrating over $u \in L$ gives

$$J_{n,R}^{(2)}(L) \leq \lambda_{(S,H)}(L) \int_{R-r_L}^\infty \phi(t-\delta)\, dV_x(t). \tag{47}$$

Combining Equations (46) and (47) with the eventual boundedness of $(a_{1,n})_n$ and $(a_{2,n})_n$, and using that a fixed δ-shift does not affect vanishing of the tail,

$$\int_0^\infty \phi(t)\, dV_x(t) < \infty,$$

we obtain

$$\lim_{R \to \infty} \sup_{n \geq N(L)} I_{n,R}(L) = 0$$

for some $N(L) \in \mathbb{N}$ with $N(L) \geq N_0$. This is exactly condition (5) of Theorem 6.4.

We now verify condition (4) of Theorem 6.4. Since

$$\Gamma_\infty = \delta_x + \eta, \qquad \eta \sim \mathrm{PPP}(c\lambda_{(S,H)}),$$

the standard Poisson moment formulas give

$$\alpha^{(2)}_{\Gamma_\infty} = c^2 \lambda^{\otimes 2}_{(S,H)} + c\,\delta_x \otimes \lambda_{(S,H)} + c\,\lambda_{(S,H)} \otimes \delta_x.$$

Fix compact $L \subset S$. Then,

$$\begin{aligned}\int_{L\times S} \overline{\kappa}_\infty(d_S(u,v))\,\alpha^{(2)}_{\Gamma_\infty}(du,dv) &= c^2 \iint_{L\times S} \overline{\kappa}_\infty(d_S(u,v))\,\lambda_{(S,H)}(du)\lambda_{(S,H)}(dv) \\ &\quad + c\,\mathbf{1}_{\{x\in L\}} \int_S \overline{\kappa}_\infty(d_S(x,v))\,\lambda_{(S,H)}(dv) \\ &\quad + c \int_L \overline{\kappa}_\infty(d_S(u,x))\,\lambda_{(S,H)}(du). \qquad (48)\end{aligned}$$

The last term is finite because $0 \leq \overline{\kappa}_\infty \leq 1$ and $\lambda_{(S,H)}(L) < \infty$. For the first two terms, Assumption 7.13 gives $\overline{\kappa}_\infty(t) \leq \phi(t)$ for $t \geq t_0$, so

$$\int_0^\infty \overline{\kappa}_\infty(t)\,dV_x(t) \leq V_x(t_0) + \int_{t_0}^\infty \phi(t)\,dV_x(t) < \infty.$$

By transitive isometric invariance, this implies

$$\int_S \overline{\kappa}_\infty(d_S(u,v))\,\lambda_{(S,H)}(dv) < \infty \qquad \text{for every } u \in S,$$

and therefore all terms in Equation (48) are finite. This is exactly condition (4) of Theorem 6.4. ◀

A.7 Lemma 7.24: Cox Radial Growth Domination

Proof. By Assumption 7.21 and Equation (34), for each $n \in \mathbb{N}$,

$$\Gamma^x_n \overset{d}{=} \delta_x + \eta_n, \qquad \eta_n \sim \mathrm{Cox}(\Phi^x_n),$$

and similarly

$$\Gamma_\infty \overset{d}{=} \delta_x + \eta, \qquad \eta \sim \mathrm{Cox}(\Phi_\infty).$$

By linearity and the monotone convergence theorem [19, Thm. 1.21], the bounds in Assumption 7.23 extend from indicators to all nonnegative compactly supported measurable f, g, h. In particular, for every nonnegative $f, g, h \in C_c(S)$ and all n,

$$\mathbb{E}\left[\int f\,d\Phi^x_n\right] \leq c_1 \int f\,d\lambda_{(S,H)},$$

$$\mathbb{E}\Big[\int f\,d\Phi_n^x \int g\,d\Phi_n^x\Big] \le c_2 \left(\int f\,d\lambda_{(S,H)}\right)\left(\int g\,d\lambda_{(S,H)}\right),$$

$$\mathbb{E}\Big[\int f\,d\Phi_n^x \int g\,d\Phi_n^x \int h\,d\Phi_n^x\Big] \le c_3 \left(\int f\,d\lambda_{(S,H)}\right)\left(\int g\,d\lambda_{(S,H)}\right)\left(\int h\,d\lambda_{(S,H)}\right).$$

Since $\Phi_n^x \Rightarrow \Phi_\infty$, the continuous mapping theorem gives convergence in distribution of these integrals for such f, g, h. The third bound gives the required uniform integrability, so the first two bounds pass to the limit. Applying Lemma 7.9, we obtain

$$\Lambda_{\eta_n}(A) \le c_1 \lambda_{(S,H)}(A), \qquad \alpha^{(2)}_{\eta_n}(A \times B) \le c_2 \lambda_{(S,H)}(A)\lambda_{(S,H)}(B),$$

and

$$\int f\,d\Lambda_\eta \le c_1 \int f\,d\lambda_{(S,H)}, \qquad \iint f(u)g(v)\,\alpha^{(2)}_\eta(du,dv) \le c_2 \left(\int f\,d\lambda_{(S,H)}\right)\left(\int g\,d\lambda_{(S,H)}\right)$$

for all nonnegative $f, g \in C_c(S)$. Hence $\Lambda_\eta \le c_1\lambda_{(S,H)}$ and $\alpha^{(2)}_\eta \le c_2 \lambda^{\otimes 2}_{(S,H)}$ as measures.

We verify condition (5) of Theorem 6.4 first. Since $\Gamma_n^x = \delta_x + \eta_n$,

$$\alpha^{(2)}_{\Gamma_n^x} = \alpha^{(2)}_{\eta_n} + \delta_x \otimes \Lambda_{\eta_n} + \Lambda_{\eta_n} \otimes \delta_x.$$

Thus the proof of Appendix A.6 applies verbatim from the decomposition of $\alpha^{(2)}_{\Gamma_n^x}$ onward, with the eventually bounded coefficients $a_{1,n}, a_{2,n}$ there replaced by the fixed constants c_1, c_2 here. This gives condition (5).

For condition (4), the same decomposition holds at the limit:

$$\alpha^{(2)}_{\Gamma_\infty} = \alpha^{(2)}_\eta + \delta_x \otimes \Lambda_\eta + \Lambda_\eta \otimes \delta_x.$$

Using the above bounds on Λ_η and $\alpha^{(2)}_\eta$, the same final integrability argument as in Appendix A.6 shows that

$$\int_{L\times S} \overline{\kappa}_\infty(d_S(u,v))\,\alpha^{(2)}_{\Gamma_\infty}(du,dv) < \infty$$

for every compact $L \subseteq S$. This is condition (4).

Hence Assumptions 7.13, 7.21 and 7.23 together give Assumption 7.5. ◀

A.8 Lemma 7.18: Count Table Verification

Proof. Write $m_n := \lambda_{(S,H)}(S_n)$ and $Y_n := N_n/m_n$. By hypothesis, $m_n \to \infty$. It is therefore enough to show, for each row of Table 3, that

$$0 < \mathbb{E}[N_n] < \infty \quad \text{for all sufficiently large } n,$$

$$Y_n \Rightarrow c, \qquad \sup_{n \geq N_0} \mathbb{E}[Y_n^3] < \infty \quad \text{for some } N_0 \in \mathbb{N}.$$

Then $N_n \to \infty$ in probability as well, since for every $M > 0$,

$$\mathbb{P}(N_n \leq M) \leq \mathbb{P}(Y_n \leq c/2) + \mathbf{1}_{\{M/m_n \geq c/2\}} \to 0.$$

Fixed. Here $N_n = \lfloor cm_n \rfloor$ almost surely. Since $m_n \to \infty$ and $c > 0$, we have $0 < \mathbb{E}[N_n] = \lfloor cm_n \rfloor < \infty$ for all sufficiently large n. Moreover,

$$Y_n = \frac{\lfloor cm_n \rfloor}{m_n} \to c$$

deterministically, and $Y_n \leq c$, so $\sup_n \mathbb{E}[Y_n^3] \leq c^3 < \infty$.

Poisson. Let $\lambda_n := \lfloor cm_n \rfloor$. Then $N_n \sim \mathrm{Poi}(\lambda_n)$, so $0 < \mathbb{E}[N_n] = \lambda_n < \infty$ for all sufficiently large n. As a standard fact,

$$\mathbb{E}[Y_n] = \frac{\lambda_n}{m_n} \to c, \qquad \mathrm{Var}(Y_n) = \frac{\lambda_n}{m_n^2} \to 0,$$

so $Y_n \to c$ in L^2, hence in distribution. Also,

$$\mathbb{E}[Y_n^3] = \frac{\lambda_n^3 + 3\lambda_n^2 + \lambda_n}{m_n^3},$$

which is uniformly bounded.

Noisy Binomial. Let $K_n := \lceil (1+\varepsilon_n) cm_n \rceil$ and $p_n := (1+\varepsilon_n)^{-1}$. Then $N_n \sim \mathrm{Bin}(K_n, p_n)$, so

$$0 < \mathbb{E}[N_n] = K_n p_n < \infty$$

for all sufficiently large n. Also,

$$\mathbb{E}[Y_n] = \frac{K_n p_n}{m_n} \to c, \qquad \mathrm{Var}(Y_n) = \frac{K_n p_n (1-p_n)}{m_n^2} \to 0,$$

so again $Y_n \to c$ in L^2. Since $N_n \leq K_n$ almost surely and K_n/m_n is bounded, $\sup_n \mathbb{E}[Y_n^3] < \infty$.

Sublinear Noise. Here

$$N_n = \lfloor cm_n + cm_n^\alpha Z_n \rfloor$$

with $0 \leq \alpha < 1$. Since $Z_n \geq 0$, we have $N_n \geq \lfloor cm_n \rfloor$, so $0 < \mathbb{E}[N_n] < \infty$ for all sufficiently large n, and finiteness is immediate from $\mathbb{E}[Z_n^3] < \infty$. Further,

$$0 \leq Y_n \leq c + cm_n^{\alpha-1} Z_n,$$

up to an $O(m_n^{-1})$ floor error. Since $\sup_n \mathbb{E}[Z_n^3] < \infty$, also $\sup_n \mathbb{E}[Z_n] < \infty$, and therefore

$$\mathbb{E}[|Y_n - c|] \leq c\, m_n^{\alpha-1} \mathbb{E}[Z_n] + m_n^{-1} \to 0.$$

Hence, $Y_n \to c$ in L^1, so in distribution. Also, for all large n,

$$Y_n^3 \le (c + cZ_n)^3,$$

since $m_n^{\alpha-1} \le 1$, and the right-hand side has uniformly bounded expectation by $\sup_n \mathbb{E}[Z_n^3] < \infty$. Thus $\sup_n \mathbb{E}[Y_n^3] < \infty$.

This concludes that every admissible count family in Table 3 satisfies Assumption 7.10. ◀

B Derivations

B.1 Derivations for Section 5.1

1) Classical Derivation.

For bounded measurable F on configurations, apply the Palm identity to

$$f((u,v),\tau) := \mathbf{1}_{\{u=v\}}\, F(\theta_{-u}\tau).$$

By stationarity and transitivity on the diagonal orbit,

$$\begin{aligned}
\mathbb{E}\left[\sum_{x\in V(\xi)} F(\theta_{-x}\xi)\right] &= \mathbb{E}\left[\sum_{(u,v)\in\xi} \mathbf{1}_{\{u=v\}}\, F(\theta_{-u}\xi)\right] \\
&= \int_{S^2} \mathbb{E}^{(u,v)}\left[\mathbf{1}_{\{u=v\}}\, F(\theta_{-u}\xi)\right] \Lambda_\xi(d(u,v)) \\
&= \int_{\Delta^2(S)} \mathbb{E}^{(x,x)}[F(\theta_{-x}\xi)]\, \Lambda_\xi(d(x,x)) \\
&= \int_{\Delta^2(S)} \mathbb{E}^{(0,0)}[F(\xi)]\, \Lambda_\xi(d(x,x)) \\
&= \Lambda_\xi(\Delta^2(S))\, \mathbb{E}^{(0,0)}[F(\xi)] \\
&= \mathbb{E}[\xi(\Delta^2(S))]\, \mathbb{E}^{(0,0)}[F(\xi)] \\
&= \mathbb{E}[|V(\xi)|]\, \mathbb{E}^{(0,0)}[F(\xi)]
\end{aligned}$$

and therefore

$$\mathbb{E}^{(0,0)}[F(\xi)] = \frac{\mathbb{E}\left[\sum_{x\in V(\xi)} F(\theta_{-x}\xi)\right]}{\mathbb{E}[|V(\xi)|]}.$$

Now by conditioning the expectation, letting $\xi_2 := \delta_{(A,A)} + \delta_{(A,B)} + \delta_{(B,A)} + \delta_{(B,B)}$, we have that

$$\begin{aligned}
\mathbb{E}\left[\sum_{x\in V(\xi)} F(\theta_{-x}\xi)\right] &= \frac{1}{2}\,\mathbb{E}[F(\delta_{(0,0)})] + \frac{1}{2}\,\mathbb{E}[F(\theta_{-A}\xi_2) + F(\theta_{-B}\xi_2)] \\
&= \frac{1}{2}\,\mathbb{E}[F(\delta_{(0,0)})] + \mathbb{E}[F(\delta_{(0,0)} + \delta_{(0,C)} + \delta_{(C,0)} + \delta_{(C,C)})].
\end{aligned}$$

Since $E[|V(\xi)|] = 3/2$, it follows that

$$\mathbb{E}^{(0,0)}[F(\xi)] = \frac{1}{3}\,\mathbb{E}[F(\delta_{(0,0)})] + \frac{2}{3}\,\mathbb{E}[F(\delta_{(0,0)} + \delta_{(0,C)} + \delta_{(C,0)} + \delta_{(C,C)})].$$

Hence, by uniqueness of probability laws under bounded measurable test functionals,

$$\xi^{(0,0)} \overset{d}{=} \begin{cases} \delta_{(0,0)}, & \text{prob. } 1/3, \\ \delta_{(0,0)} + \delta_{(C,C)} + \delta_{(0,C)} + \delta_{(C,0)}, & \text{prob. } 2/3. \end{cases}$$

2) Vertex-Normalized Derivation.

For bounded measurable F on configurations, apply Definition 19 with

$$f((u,v),\tau) := \mathbf{1}_{\{u=v\}}\, F(\theta_{-u}\tau).$$

By stationarity and transitivity on the diagonal orbit,

$$\begin{aligned}
\mathbb{E}\left[\sum_{x\in V(\xi)} \frac{1}{|V(\xi)|} F(\theta_{-x}\xi)\right] &= \mathbb{E}\left[\sum_{(u,v)\in\xi} \frac{\mathbf{1}_{\{u=v\}}}{|V(\xi)|}\, F(\theta_{-u}\xi)\right] \\
&= \int_{S^2} \hat{\mathbb{E}}^{(u,v)}\left[\mathbf{1}_{\{u=v\}}\, F(\theta_{-u}\xi)\right] \mathbb{E}^{(u,v)}\left[\frac{1}{|V(\xi)|}\right] \Lambda_\xi(d(u,v)) \\
&= \int_{\Delta^2(S)} \hat{\mathbb{E}}^{(x,x)}[F(\theta_{-x}\xi)]\, \mathbb{E}^{(x,x)}\left[\frac{1}{|V(\xi)|}\right] \Lambda_\xi(d(x,x)) \\
&= \Lambda_\xi(\Delta^2(S))\, \mathbb{E}^{(0,0)}\left[\frac{1}{|V(\xi)|}\right]\, \hat{\mathbb{E}}^{(0,0)}[F(\xi)] \\
&= \mathbb{E}[|V(\xi)|]\, \mathbb{E}^{(0,0)}\left[\frac{1}{|V(\xi)|}\right]\, \hat{\mathbb{E}}^{(0,0)}[F(\xi)].
\end{aligned}$$

As established, $\mathbb{E}[|V(\xi)|] = 3/2$ and

$$\mathbb{E}^{(0,0)}\left[\frac{1}{|V(\xi)|}\right] = \frac{1}{3}\cdot 1 + \frac{2}{3}\cdot\frac{1}{2} = \frac{2}{3},$$

so the prefactor is 1. Therefore,

$$\hat{\mathbb{E}}^{(0,0)}[F(\xi)] = \mathbb{E}\left[\sum_{x\in V(\xi)} \frac{1}{|V(\xi)|} F(\theta_{-x}\xi)\right].$$

Now condition this expectation exactly as above (with $\xi_2 = \delta_{(A,A)} + \delta_{(A,B)} + \delta_{(B,A)} + \delta_{(B,B)}$):

$$\begin{aligned}
\hat{\mathbb{E}}^{(0,0)}[F(\xi)] &= \frac{1}{2}\,\mathbb{E}[F(\delta_{(0,0)})] + \frac{1}{2}\,\mathbb{E}\left[\frac{1}{2}\big(F(\theta_{-A}\xi_2) + F(\theta_{-B}\xi_2)\big)\right] \\
&= \frac{1}{2}\,\mathbb{E}[F(\delta_{(0,0)})] + \frac{1}{2}\,\mathbb{E}\big[F\big(\delta_{(0,0)} + \delta_{(0,C)} + \delta_{(C,0)} + \delta_{(C,C)}\big)\big].
\end{aligned}$$

Hence, by uniqueness of probability laws under bounded measurable test functionals,

$$\hat{\xi}^{(0,0)} \overset{d}{=} \begin{cases} \delta_{(0,0)}, & \text{prob. } 1/2, \\ \delta_{(0,0)} + \delta_{(C,C)} + \delta_{(0,C)} + \delta_{(C,0)}, & \text{prob. } 1/2. \end{cases}$$

B.2 Detail on Table 1 Entries

Here we verify that all of the specified connection functions decay suitably, and specify the parameters under which this holds. Specifically, we give the parameters under which

$$\sup_{n \geq N_0} \overline{\kappa}_n(t) \leq Ct^{-\alpha}, \qquad \alpha > d$$

for all $t \geq t_0$, under suitable $N_0 \in \mathbb{N}$, $t_0 > 0$, and $C > 0$, where d is the parameter from Assumption 7.16. For rows whose table entry does not specify a canonical finite sequence, we take

$$\kappa_n := \kappa_\infty \qquad \text{for all } n \in \mathbb{N},$$

so that the required uniform-in-n tail bound reduces directly to a bound on $\overline{\kappa}_\infty$. For THRG and PHRG, the transformed distance coordinate is one-dimensional, so those rows are taken only with the corresponding parameter choice $d = 1$.

PSIRG

This model does not have a canonical mark space or mark distribution. This model constructs κ_∞ (before truncation to $[0, 1]$) via

$$\kappa_\infty(t, w_1, w_2) := f(t)g(w_1, w_2).$$

Here, f maps $\mathbb{R}^+ \to \mathbb{R}^+$ and g maps $\mathbb{M}^2 \to \mathbb{R}^+$. For this row, take $\kappa_n := \kappa_\infty$ for every $n \in \mathbb{N}$.

We require that f satisfies

$$f(t) \leq C_1 t^{-\beta_1}$$

for all $t \geq t_1$, under some $C_1 > 0$, $\beta_1 > 1$, and $t_1 \in \mathbb{R}^+$. We require that

$$\mathbb{P}(g(W_1, W_2) > t) \leq C_2 t^{-\beta_2}$$

for all $t \geq t_2$, under some $C_2 > 0$, $\beta_2 > 0$, and $t_2 \in \mathbb{R}^+$, where $W_1, W_2 \overset{i.i.d.}{\sim} \nu$. We further require that

$$\min(\beta_1, \beta_1\beta_2) > d.$$

Choose $\varepsilon > 0$ small enough that

$$\alpha := \min(\beta_1, \beta_1\beta_2) - \varepsilon > d.$$

These are all the required conditions. Indeed, by [18, Lem. 2.3], it holds that

$$\overline{\kappa}_\infty(t) \leq C_3 t^{-\alpha}$$

for all sufficiently large t, under some $C_3 > 0$. Since $\kappa_n = \kappa_\infty$ for every n, we also have $\overline{\kappa}_n = \overline{\kappa}_\infty$ for every n. Thus, this model satisfies Assumption 7.16 with parameter d under the stated specification.

GIRG$_1$

This is the $\alpha_G = \infty$ case of [18, Sec. 2.1.2, Cor. 2.8, and proof]. This model has a canonical mark space $\mathbb{M} = [1, \infty)$ and canonical mark distribution ν satisfying

$$\nu([1, w]) = 1 - w^{1-\beta}$$

for all $w \geq 1$, under some $\beta > 2$. This model constructs κ_∞ via

$$\kappa_\infty(t, w_1, w_2) := \mathbf{1}\left\{\left(\frac{w_1 w_2}{\mathbb{E}[W]}\right)^{1/d} > t\right\},$$

where $W \sim \nu$. For this row, take $\kappa_n := \kappa_\infty$ for every $n \in \mathbb{N}$.

We require that the mark law satisfies

$$c_G z^{1-\beta} \leq \nu([z, \infty]) \leq C_G z^{1-\beta}$$

for all $z \geq z_0$, under some $c_G, C_G, z_0 > 0$ and $\beta > 2$.

This is the only required condition. Indeed, fix any $\gamma > d$. Then

$$\kappa_\infty(t, w_1, w_2) \leq \left(\frac{w_1 w_2}{\mathbb{E}[W]}\right)^{\gamma/d} t^{-\gamma}.$$

Thus, this model is dominated by a PSIRG kernel with

$$f(t) := t^{-\gamma}, \qquad g(w_1, w_2) := \left(\frac{w_1 w_2}{\mathbb{E}[W]}\right)^{\gamma/d}.$$

By Breiman's lemma exactly as used in the cited proof, for every sufficiently small $\varepsilon > 0$, there exists $C_\varepsilon > 0$ such that

$$\mathbb{P}(g(W_1, W_2) > t) \leq C_\varepsilon t^{-\beta_2}$$

for all sufficiently large t, where

$$\beta_2 := \frac{d(\beta - 1 - \varepsilon)}{\gamma}.$$

Since $\beta > 2$, one can choose $\varepsilon > 0$ small enough that $\beta - 1 - \varepsilon > 1$, and then

$$\alpha_* := \min(\gamma, \gamma\beta_2) = \min(\gamma, d(\beta - 1 - \varepsilon)) > d.$$

Choose $\delta > 0$ small enough that $\alpha := \alpha_* - \delta > d$. Therefore, by the PSIRG bound from [18, Lem. 2.3], it holds that

$$\overline{\kappa}_\infty(t) \leq C t^{-\alpha}$$

for all sufficiently large t, under some $C > 0$. Since $\kappa_n = \kappa_\infty$ for every n, we also have $\overline{\kappa}_n = \overline{\kappa}_\infty$ for every n. Thus, this model satisfies Assumption 7.16 with parameter d under the stated specification.

GIRG$_2$

This is the $1 < \alpha_G < \infty$ case of [18, Sec. 2.1.2, Cor. 2.8, and proof]. This model has a canonical mark space $\mathbb{M} = [1, \infty)$ and canonical mark distribution ν satisfying

$$\nu([1, w]) = 1 - w^{1-\beta}$$

for all $w \geq 1$, under some $\beta > 2$. This model constructs κ_∞ (before truncation to $[0, 1]$) via

$$\kappa_\infty(t, w_1, w_2) := \left(\frac{w_1 w_2}{\mathbb{E}[W]}\right)^{\alpha_G} t^{-d\alpha_G},$$

where $W \sim \nu$. For this row, take $\kappa_n := \kappa_\infty$ for every $n \in \mathbb{N}$.

We require that the mark law satisfies

$$c_G z^{1-\beta} \leq \nu([z, \infty)) \leq C_G z^{1-\beta}$$

for all $z \geq z_0$, under some $c_G, C_G, z_0 > 0$.

This is the only required condition. Indeed, this model is already of PSIRG form, with

$$f(t) := t^{-d\alpha_G}, \qquad g(w_1, w_2) := \left(\frac{w_1 w_2}{\mathbb{E}[W]}\right)^{\alpha_G}.$$

By Breiman's lemma exactly as used in the cited proof, for every sufficiently small $\varepsilon > 0$, there exists $C_\varepsilon > 0$ such that

$$\mathbb{P}(g(W_1, W_2) > t) \leq C_\varepsilon t^{-\beta_2}$$

for all sufficiently large t, where

$$\beta_2 := \frac{\beta - 1 - \varepsilon}{\alpha_G}.$$

Since $\alpha_G > 1$ and $\beta > 2$, one can choose $\varepsilon > 0$ small enough that $\beta - 1 - \varepsilon > 1$, and then

$$\alpha_* := \min(d\alpha_G, d\alpha_G\beta_2) = \min(d\alpha_G, d(\beta - 1 - \varepsilon)) > d.$$

Choose $\delta > 0$ small enough that $\alpha := \alpha_* - \delta > d$. Therefore, by [18, Lem. 2.3], it holds that

$$\overline{\kappa}_\infty(t) \leq Ct^{-\alpha}$$

for all sufficiently large t, under some $C > 0$. Since $\kappa_n = \kappa_\infty$ for every n, we also have $\overline{\kappa}_n = \overline{\kappa}_\infty$ for every n. Thus, this model satisfies Assumption 7.16 with parameter d under the stated specification.

THRG

This model has a canonical mark space $\mathbb{M} = [1, \infty)$, canonical mark distribution ν satisfying

$$\nu([1, w]) = 1 - w^{-2\alpha_H}$$

for all $w \geq 1$, under some $\alpha_H > 1/2$. This model has a canonical converging sequence given by

$$\kappa_n(t, w_1, w_2) := \mathbf{1}\left\{D_n(t, w_1, w_2) \leq R_n\right\},$$

and limiting connection function

$$\kappa_\infty(t, w_1, w_2) := \mathbf{1}\left\{t \leq \frac{\beta w_1 w_2}{\pi}\right\}.$$

Here, $\beta > 0$ is the table parameter,

$$R_n := 2\log(n/\beta),$$

and

$$r_{n,1} := R_n - 2\log w_1, \qquad r_{n,2} := R_n - 2\log w_2.$$

Also, $D_n(t, w_1, w_2)$ denotes the hyperbolic distance between two points with angular separation $2\pi t/n$ and radial coordinates $r_{n,1}, r_{n,2}$, namely

$$D_n(t, w_1, w_2) := \operatorname{arcosh}\left(\cosh r_{n,1} \cosh r_{n,2} - \sinh r_{n,1} \sinh r_{n,2} \cos\left(\frac{2\pi t}{n}\right)\right).$$

This distance mimics the hyperbolic distance via the real distance (representing angular distance) and marks (representing radial positions). In particular, the relevant polynomial ball-growth exponent for this transformed model is $d = 1$.

This is the only required condition. Indeed, by [18, Sec. 2.1.3, Cor. 2.9, and proof], fix any $\gamma > 1$. Then

$$\kappa_\infty(t, w_1, w_2) \leq \left(\frac{\beta w_1 w_2}{\pi}\right)^\gamma t^{-\gamma}.$$

It also remains to verify sequential convergence of the canonical finite sequence to the stated limit. Fix any convergent sequence

$$(t_n, w_{1,n}, w_{2,n}) \to (t, w_1, w_2)$$

in $\mathbb{R}^+ \times [1, \infty)^2$, where (t, w_1, w_2) is a continuity point of κ_∞. Write $r_{n,i} := R_n - 2\log w_{i,n}$. Using

$$\cosh D_n = \cosh(r_{n,1} - r_{n,2}) + 2\sinh(r_{n,1})\sinh(r_{n,2})\sin^2\left(\frac{\pi t_n}{n}\right),$$

we have $r_{n,i} \to \infty$, so $e^{-R_n}\cosh(r_{n,1}-r_{n,2}) \to 0$. Also,

$$4e^{-R_n}\sinh(r_{n,1})\sinh(r_{n,2})\sin^2\left(\frac{\pi t_n}{n}\right) \to \left(\frac{\pi t}{\beta w_1 w_2}\right)^2,$$

since $e^{r_{n,1}+r_{n,2}-R_n} = e^{R_n}/(w_{1,n}^2 w_{2,n}^2) = n^2/(\beta^2 w_{1,n}^2 w_{2,n}^2)$ and $n\sin(\pi t_n/n) \to \pi t$. Hence, if $t > 0$, then

$$e^{(D_n(t_n,w_{1,n},w_{2,n})-R_n)/2} \to \frac{\pi t}{\beta w_1 w_2},$$

using $\operatorname{arcosh}(y) = \log(2y) + o(1)$ as $y \to \infty$. If $t = 0$, then the same computation gives $e^{-R_n}\cosh D_n \to 0$, so $D_n - R_n \to -\infty$, and therefore the same limit holds with value 0. Thus

$$\mathbf{1}\{D_n(t_n, w_{1,n}, w_{2,n}) \le R_n\} \to \mathbf{1}\left\{\frac{\pi t}{\beta w_1 w_2} \le 1\right\} = \mathbf{1}\left\{t \le \frac{\beta w_1 w_2}{\pi}\right\},$$

and since (t, w_1, w_2) is a continuity point of κ_∞, equality does not occur in the limit. Therefore $\kappa_n \xrightarrow{s} \kappa_\infty$.

The cited proof only gave a bound for the limiting kernel, so we still need a direct bound for the canonical finite sequence. If $\kappa_n(t, w_1, w_2) = 1$, then $D_n(t, w_1, w_2) \le R_n$. Since t is the one-dimensional torus distance coordinate, we have $0 \le t \le n/2$. If $r_{n,1}, r_{n,2} \ge 1$, then the hyperbolic law of cosines gives

$$2\sinh(r_{n,1})\sinh(r_{n,2})\sin^2\left(\frac{\pi t}{n}\right) \le \cosh(R_n).$$

Using $\sinh u \ge e^u/4$ for $u \ge 1$, $\cosh(R_n) \le e^{R_n}/2$, and $\sin x \ge 2x/\pi$ on $[0, \pi/2]$, it follows that

$$\frac{e^{r_{n,1}+r_{n,2}}}{8}\,\frac{4t^2}{n^2} \le \frac{e^{R_n}}{2},$$

hence

$$t \le \beta w_1 w_2.$$

If instead $r_{n,1} < 1$ or $r_{n,2} < 1$, then $w_1 > e^{(R_n-1)/2}$ or $w_2 > e^{(R_n-1)/2}$, so

$$\beta w_1 w_2 \ge \beta e^{(R_n-1)/2} = e^{-1/2} n > \frac{n}{2},$$

and therefore again $t \le \beta w_1 w_2$. Thus

$$\kappa_n(t, w_1, w_2) \le \mathbf{1}\{t \le \beta w_1 w_2\} \le (\beta w_1 w_2)^\gamma t^{-\gamma}.$$

Thus, the canonical finite sequence is dominated by a PSIRG kernel with

$$f(t) := t^{-\gamma}, \qquad g(w_1, w_2) := (\beta w_1 w_2)^\gamma.$$

By Breiman's lemma exactly as used in the cited proof, for every sufficiently small $\varepsilon > 0$, there exists $C_\varepsilon > 0$ such that

$$\mathbb{P}(g(W_1, W_2) > t) \le C_\varepsilon t^{-\beta_2}$$

for all sufficiently large t, where

$$\beta_2 := \frac{2\alpha_H - \varepsilon}{\gamma}.$$

Since $\alpha_H > 1/2$, one can choose $\varepsilon > 0$ small enough that $2\alpha_H - \varepsilon > 1$, and then

$$\alpha_* := \min(\gamma, \gamma\beta_2) = \min(\gamma, 2\alpha_H - \varepsilon) > 1 = d.$$

Choose $\delta > 0$ small enough that $\alpha := \alpha_* - \delta > 1 = d$. Therefore, by the PSIRG bound [18, Lem. 2.3], it holds that

$$\overline{\kappa}_n(t) \leq Ct^{-\alpha}$$

eventually, under some $C > 0$. Thus, this model satisfies Assumption 7.16 with parameter $d = 1$ under the stated specification.

PHRG

This model has a canonical mark space $\mathbb{M} = [1, \infty)$, canonical mark distribution ν satisfying

$$\nu([1, w]) = 1 - w^{-2\alpha_H}$$

for all $w \geq 1$, under some $\alpha_H > 1/2$. This model has a canonical converging sequence given by

$$\kappa_n(t, w_1, w_2) := \left(1 + \exp\left(\frac{D_n(t, w_1, w_2) - R_n}{2T}\right)\right)^{-1},$$

and limiting connection function

$$\kappa_\infty(t, w_1, w_2) := \left(1 + \left(\frac{\pi t}{\beta w_1 w_2}\right)^{1/T}\right)^{-1},$$

where

$$R_n := 2\log(n/\beta),$$

$\beta > 0$ is the table parameter, $0 < T < 1$, and D_n is the same hyperbolic distance function as above. In particular, the relevant polynomial ball-growth exponent for this transformed model is again $d = 1$.

These are the only required conditions. Indeed, by [18, Sec. 2.1.3, Cor. 2.9, and proof], using that $(1 + u)^{-1} \leq 1 \wedge u^{-1}$ for all $u \geq 0$, it follows that

$$\kappa_\infty(t, w_1, w_2) \leq 1 \wedge \left(\frac{\beta w_1 w_2}{\pi}\right)^{1/T} t^{-1/T}.$$

It also remains to verify sequential convergence of the canonical finite sequence. Fix any convergent sequence

$$(t_n, w_{1,n}, w_{2,n}) \to (t, w_1, w_2)$$

in $\mathbb{R}^+ \times [1,\infty)^2$. With $r_{n,i} := R_n - 2\log w_{i,n}$, the same computation as in the THRG case gives

$$e^{(D_n(t_n,w_{1,n},w_{2,n})-R_n)/2} \to \frac{\pi t}{\beta w_1 w_2},$$

where the limit is interpreted as 0 when $t = 0$. Therefore,

$$\exp\left(\frac{D_n(t_n,w_{1,n},w_{2,n}) - R_n}{2T}\right) \to \left(\frac{\pi t}{\beta w_1 w_2}\right)^{1/T},$$

and hence, by continuity of $u \mapsto (1+u)^{-1}$,

$$\kappa_n(t_n,w_{1,n},w_{2,n}) \to \left(1 + \left(\frac{\pi t}{\beta w_1 w_2}\right)^{1/T}\right)^{-1} = \kappa_\infty(t,w_1,w_2).$$

Thus $\kappa_n \xrightarrow{s} \kappa_\infty$.

The cited proof again only gives a bound for the limiting kernel, so we still need a direct bound for the canonical finite sequence. Since $(1+e^x)^{-1} \le 1 \wedge e^{-x}$ for all $x \in \mathbb{R}$, we have

$$\kappa_n(t,w_1,w_2) \le 1 \wedge \exp\left(-\frac{D_n(t,w_1,w_2) - R_n}{2T}\right).$$

If $r_{n,1}, r_{n,2} \ge 1$, then the hyperbolic law of cosines, together with $\sinh u \ge e^u/4$ for $u \ge 1$ and $\sin x \ge 2x/\pi$ on $[0,\pi/2]$, gives

$$\cosh(D_n(t,w_1,w_2)) \ge 2\sinh(r_{n,1})\sinh(r_{n,2})\sin^2\left(\frac{\pi t}{n}\right) \ge \frac{n^2 t^2}{2\beta^4 w_1^2 w_2^2}.$$

Since $\cosh u \le e^u$ for all $u \ge 0$, it follows that

$$D_n(t,w_1,w_2) - R_n \ge 2\log\left(\frac{t}{\sqrt{2}\,\beta w_1 w_2}\right).$$

Hence

$$\kappa_n(t,w_1,w_2) \le 1 \wedge 2^{1/(2T)}\left(\frac{\beta w_1 w_2}{t}\right)^{1/T}.$$

If instead $r_{n,1} < 1$ or $r_{n,2} < 1$, then as above $\beta w_1 w_2 > n/2 \ge t$, so the same upper bound still holds after enlarging the constant if necessary. Thus, the canonical finite sequence is dominated by a PSIRG kernel with

$$f(t) := t^{-1/T}, \qquad g(w_1,w_2) := 2^{1/(2T)}(\beta w_1 w_2)^{1/T}.$$

By Breiman's lemma exactly as used in the cited proof, for every sufficiently small $\varepsilon > 0$, there exists $C_\varepsilon > 0$ such that

$$\mathbb{P}(g(W_1,W_2) > t) \le C_\varepsilon t^{-\beta_2}$$

for all sufficiently large t, where

$$\beta_2 := (2\alpha_H - \varepsilon)T.$$

Since $0 < T < 1$ and $\alpha_H > 1/2$, one can choose $\varepsilon > 0$ small enough that $2\alpha_H - \varepsilon > 1$, and then

$$\alpha_* := \min(1/T, (1/T)\beta_2) = \min(1/T, 2\alpha_H - \varepsilon) > 1 = d.$$

Choose $\delta > 0$ small enough that $\alpha := \alpha_* - \delta > 1 = d$. Therefore, by the PSIRG bound [18, Lem. 2.3], it holds that

$$\overline{\kappa}_n(t) \le Ct^{-\alpha}$$

eventually, under some $C > 0$. Thus, this model satisfies Assumption 7.16 with parameter $d = 1$ under the stated specification.

CSFP - λ, α

This model has a canonical mark space $\mathbb{M} = [1, \infty)$ and canonical mark distribution ν satisfying

$$\nu([1, w]) = 1 - w^{-\beta}$$

for all $w \ge 1$, under some $\beta > 0$. This model constructs κ_∞ via

$$\kappa_\infty(t, w_1, w_2) := 1 - \exp(-\lambda w_1 w_2 t^{-\alpha}),$$

under some $\lambda > 0$. For this row, take $\kappa_n := \kappa_\infty$ for every $n \in \mathbb{N}$.

These are the only required conditions. Indeed, by [18, Sec. 2.1.4, Cor. 2.10, and proof] and using that $1 - e^{-x} \le x$ for all $x \ge 0$, it follows that

$$\kappa_\infty(t, w_1, w_2) \le \lambda w_1 w_2 t^{-\alpha}.$$

Thus, this model is dominated by a PSIRG kernel with

$$f(t) := t^{-\alpha}, \qquad g(w_1, w_2) := \lambda w_1 w_2.$$

By Breiman's lemma exactly as used in the cited proof, for every sufficiently small $\varepsilon > 0$, there exists $C_\varepsilon > 0$ such that

$$\mathbb{P}(g(W_1, W_2) > t) \le C_\varepsilon t^{-(\beta - \varepsilon)}$$

for all sufficiently large t. Hence

$$\alpha' := \min(\alpha, \alpha(\beta - \varepsilon)).$$

If

$$\min(\alpha, \alpha\beta) > d,$$

then one can choose $\varepsilon > 0$ small enough that $\alpha' > d$. Therefore, by [18, proof of Cor. 2.10], it holds that

$$\overline{\kappa}_\infty(t) \le Ct^{-\alpha'}$$

for all sufficiently large t, under some $C > 0$. Since $\kappa_n = \kappa_\infty$ for every n, we also have $\overline{\kappa}_n = \overline{\kappa}_\infty$ for every n. Thus, this model satisfies Assumption 7.16 with parameter d under the stated specification.

WDRCM - ρ, h

This model has a canonical mark space $\mathbb{M} = [0,1]$ and canonical mark distribution ν satisfying

$$\nu([0,w]) = w$$

for all $w \in [0,1]$. This model constructs κ_∞ via

$$\kappa_\infty(t, w_1, w_2) := \rho(h(w_1, w_2, t)),$$

where $\rho : \mathbb{R}^+ \to [0,1]$ and $h : [0,1]^2 \times \mathbb{R}^+ \to \mathbb{R}^+$ are measurable, as in [18, Sec. 2.1.5]. For this row, take $\kappa_n := \kappa_\infty$ for every $n \in \mathbb{N}$.

We require that there exist $C > 0$, $\alpha > d$, and $t_0 \in \mathbb{R}^+$ such that

$$\overline{\kappa}_\infty(t) = \mathbb{E}[\rho(h(W_1, W_2, t))] \leq Ct^{-\alpha}$$

for all $t \geq t_0$, where $W_1, W_2 \overset{i.i.d.}{\sim} \nu$.

This is the only required condition. Indeed,

$$\overline{\kappa}_n(t) = \overline{\kappa}_\infty(t) \leq Ct^{-\alpha}$$

for all $t \geq t_0$ and every $n \in \mathbb{N}$. Thus, this model satisfies Assumption 7.16 with parameter d under the stated specification.

SDPARG - c, α

This model is the shell-decay p-adic random graph in the spirit of [8]. This model constructs κ_∞ via

$$\kappa_\infty(t) := c(t)t^{-\alpha},$$

where $c : \mathbb{R}^+ \to \mathbb{R}^+$ is measurable and bounded, and where the model is markless. For this row, take $\kappa_n := \kappa_\infty$ for every $n \in \mathbb{N}$.

These are the only required conditions. Indeed, let

$$C_0 := \sup_{t \geq 0} c(t) < \infty.$$

Then

$$\overline{\kappa}_\infty(t) = \kappa_\infty(t) \leq C_0 t^{-\alpha}$$

for all $t > 0$. If $\alpha > d$, then $\overline{\kappa}_n = \overline{\kappa}_\infty$ for every n, so

$$\overline{\kappa}_n(t) \leq C_0 t^{-\alpha}$$

for all $t > 0$ and every $n \in \mathbb{N}$. Thus, this model satisfies Assumption 7.16 with parameter d under the stated specification.

C Glossary of Notation

Notation	Meaning	Reference
S	LC Polish space	Def. 1
d_S	Fixed proper compatible metric on S	Def. 1
S^2	Square LC Polish space	Def. 1
τ	Adjacency measure	Def. 3
μ, ν	Counting measure	Def. 3
$\mathcal{N}(S^2)$	Space of boundedly finite simple counting measures on S^2	Def. 3
$\mathcal{A}(S^2)$	Space of adjacency measures on S^2	Def. 3
(τ, x)	Rooted adjacency measure	Def. 4
$\mathcal{A}^\circ(S^2)$	Space of rooted adjacency measures on S^2	Def. 4
ξ	Adjacency process	Def. 5
(ξ, p)	Rooted adjacency process	Def. 5
η	Point process	Def. 5
δ_z	The measure taking $A \mapsto \mathbf{1}_{z \in A}$	Def. 6
$\mathcal{A}_{\mathrm{sym}}(S^2)$	Space of symmetric adjacency measures on S^2	Def. 6
$\mathcal{A}^\circ_{\mathrm{sym}}(S^2)$	Space of rooted symmetric adjacency measures on S^2	Def. 6
$\mathrm{Sym}(\tau)$	Symmetrization map on adjacency measures	Def. 6
$\mathrm{Sym}(\tau, x)$	Rooted symmetrization map (lift of Sym)	Def. 6
V	Vertex map	Def. 7
$\mathcal{G}_*$	Benjamini–Schramm space of rooted locally finite graphs	Sec. 3.2
$G_S(\tau)$	Constructed labeled graph from τ	Def. 8
$G_S(\tau, x)$	Rooted constructed labeled graph from (τ, x)	Def. 8
V_τ	Vertex set from τ	Def. 8
E_τ	Edge set from τ	Def. 8
$[G]$, $[G, x]$	(Rooted) graph isomorphism class	Def. 9
$G(\tau)$	Constructed graph	Def. 9
$G(\tau, x)$	Local rooted constructed graph	Def. 9
$\mathcal{C}_{G_S(\tau)}(x)$	Connected component of x in $G_S(\tau)$	Def. 9
$\mathcal{E}(S^2)$	Space of regular adjacency measures	Def. 11
$\mathcal{E}^\circ(S^2)$	Space of rooted regular adjacency measures	Def. 11
B_r	Rooted radius-r ball map on rooted graphs	Def. 12
V_r	Radius-r vertex-counting map	Def. 12
$d_G(x, y)$	Graph distance in a graph G	Def. 12
R	Rooting kernel on a subspace $\mathcal{M} \subseteq \mathcal{E}(S^2)$	Def. 13
$\mathbb{U}$	Uniform rooting kernel	Def. 14
K_n	Complete graph on n vertices	Sec. 4

$\tau\|_X$	Restriction of an adjacency measure to a set $X \subseteq S^2$	Eq. (1)
H	A unimodular LCSC topological group	Eq. (1)
θ_h	Shift operator $\mu \mapsto \sum_{z\in\mu} \delta_{h\cdot z}$	Eq. (7)
(S, H)	Grouped LC Polish space	Def. 16
$(S, H)^2$	Square grouped LC Polish space	Def. 17
Λ_η	Intensity measure of η	Eq. (9)
$\mathbb{P}^x$	Palm kernel at $x \in S$	Def. 18
η^x	Palm version of η at x	Def. 18
$\xi^{(x,x)}$	Palm version of adjacency process ξ at (x, x)	Eq. (11)
$\Delta^2(X)$	Diagonal image of a set X	Def. 19
$\hat{\mathbb{P}}^z$	Vertex-normalized Palm kernel at $z \in S^2$	Def. 19
$\hat{\xi}^z$	Vertex-normalized Palm version of ξ at z	Def. 19
$\subset_{\alpha,x}$	α-admissible sequence sharing (x, x)	Def. 20
$\mathbb{T}_n^d$	Volume-n flat d-torus	Def. 20
$\|\cdot\|_{\mathrm{TV}}$	Total variation norm	Eq. (12)
C	Connection kernel	Def. 21
$\mathbb{M}$	Mark space for SIRGs	Def. 22
M	Marking kernel	Def. 22
κ	Connection function	Def. 22
Q_κ	Edge-construction kernel on marked configurations	Def. 22
$\subset_{\beta,x}$	β-admissible sequence sharing (x, x)	Def. 23
$\xrightarrow{s}$	Sequential convergence (at continuity points of the limit)	Sec. 6.3
$\overset{s}{\Longrightarrow}$	Sequential weak convergence of kernels	Sec. 6.4
$\alpha_\eta^{(2)}$	Second factorial moment measure of η	Eq. (27)
$\overline{\kappa}^\nu$	Mark-averaged connection function under mark law ν	Sec. 6.4
$\overline{\kappa}$	Shorthand for $\overline{\kappa}^\nu$ when ν is clear	Sec. 6.4
$D(\eta)$	Distance process associated with η	Sec. 6.4
$\pi_k(x_1, ..., x_n)$	Projection to x_k	Prop. 6.5
N^*	Size-biased version of N	Eq. (32)
$\eta^{!x}$	Reduced Palm version of η at x	Sec. 7.2
Φ	Directing measure	Sec. 7.2
$\mathcal{M}(S)$	Space of boundedly finite Borel measures on S	Sec. 7.3
Φ^x	Palm version of directing measure Φ at x	Eq. (33)
ϕ	Tail envelope used to dominate $(\overline{\kappa}_n)_n$	Prop. 7.15
$V_x(t)$	Radial volume profile	Prop. 7.15
d_p	p-adic metric, $d_p(a, b) = \|a - b\|_p$	Sec. 7.3
$\|\cdot\|_p$	p-adic absolute value (norm)	Sec. 7.3
$\mathbb{Q}_p$	Field of p-adic numbers	Sec. 7.3
$\mathbb{Z}_p$	Ring of p-adic integers (unit ball in $\mathbb{Q}_p$)	Sec. 7.3

λ_p	Normalized Haar measure on $(\mathbb{Q}_p, +)$	Sec. 7.3
q_n, q_∞	Periodic profile functions	Assumps. 7.26 and 7.29
Θ	Uniform phase shift on $[-1/2, 1/2)^d$	Eq. (42)
$C_{r,H}$	Rooted-ball cylinder set in $\mathcal{G}_*$	Appendix A.3

D Glossary of Shorthand

Notation	Short For	Reference
$z \in \mu$	$z \in \{x : \mu(\{x\}) > 0\}$	Def. 7
$\lvert\mu\rvert$	$\lvert\{x : \mu(\{x\}) > 0\}\rvert$	Def. 7
$A \cdot \mu$	$A \cdot \{x : \mu(\{x\}) > 0\}$ $\quad \cdot = \cup, \cap, \times, \ldots$	Def. 7
$B_r(\tau, x)$	$B_r(G(\tau, x))$	Def. 12
$\Rightarrow$	Convergence in distribution	Sec. 4
$=^d$, $\stackrel{d}{=}$	Equality in distribution	Sec. 4
$\asymp$	Two-sided asymptotic bound	Sec. 7.3

E Glossary of Citations

Reference	Detail
[1, Lem. 1.3]	Sequential mapping theorem
[2, Prop. 2.8]	Atom matching for counting measures
[2, Prop. 2.8]	Counting measure restriction continuous
[2, Sec. 4, p. 5]	Laplace functionals
[5, Thm. 6.7.4]	Borel embedding of Polish space into $\mathbb{R}$
[11, Prop. 2.1, Eq. (2.5)]	Tail control for edges longer than a cutoff
[13, p. 462]	Palm kernel uniqueness
[17, Def. A.3]	Rooted-ball map is measurable in Benjamini-Schramm topology
[17, Eq. (1.4.18)]	Size biased counts exist
[18, Sec. 2.1.4, Cor. 2.10, and proof]	CSFP local limit
[18, Sec. 2.1.5]	WDRCM local limit
[18, proof of Cor. 2.10]	CSFP local limit
[18, Def. 1.2]	Cylinder sets generate Benjamini-Schramm space
[18, Sec. 2.1.2, Cor. 2.8, and proof]	GIRG local limit
[18, Sec. 2.1.3, Cor. 2.9, and proof]	HRG local limit
[18, Lem. 2.3]	PSIRG Bound
[18, Thm. 1.9]	Original SIRG convergence theorem
[18, Def. 1.3]	Weak convergence as isomorphism probabilities
[19, Thm. 1.21]	Monotone convergence theorem
[19, Thm. 23.21]	Cox convergence follows from convergence of directing measures
[19, Thm. 30.5]	Convergence of Mixed Binomial to Cox processes
[19, Thm. 8.5]	Regular conditional law
[22, Def. 8.25]	Markov kernel
[22, Def. 13.26]	Prokhorov tightness
[22, Thm. 13.6]	Random variables on Polish spaces are tight
[22, Thm. 5.3(v)]	Triangle inequality for expectations
[23, Eq. 2.1]	Haar measure
[23, Eq. 2.4]	Measurability of shift operator
[24, Prop. 3.13]	Shift identity for stationary processes
[25, Thm. 4.1, Eq. (4.2)]	Mecke's formula
[26, Eq. (4.8)]	Palm and weighted Palm relation
[29, Def. 4.2.1]	Homeomorphism definition
[30, p. 28]	Vague topology

[30, p. 28]	Vauge space is Polish
[30, p. 28]	Vauge topology generated by evaluation maps
[32, p.725]	Stationarity of Cox process under stationary directing measure
[32, p. 725]	Stationarity of Cox process under stationary directing measure
[34, Def. 31]	Composition of Markov kernels
[35, Prop. 1.4]	Closed subsets of Polish space are Polish
[35, Sec. 17]	Group actions examples
[36, Def. 4.5]	Joint exchangeability
[36, Def. 4.1]	Veitch-Roy adjacency measure
[37, Thm. 8]	Palm representation of stationary Cox process
[38, Prop 18.2]	Asymptotic TV equality gives weak convergence

References

[1] Erik Avelin. *The Low-Density Lorentz Gas with Non-Identical Scatterers*. Tech. rep. U.U.D.M. Project Report 2024:11. Uppsala University, 2024.

[2] Bojan Basrak and Hrvoje Planinić. "A note on vague convergence of measures". In: *Statistics & Probability Letters* 153 (Oct. 2019), 180–186. DOI: 10.1016/j.spl.2019.06.004. URL: http://dx.doi.org/10.1016/j.spl.2019.06.004.

[3] Itai Benjamini and Oded Schramm. *Recurrence of Distributional Limits of Finite Planar Graphs*. 2001. arXiv: math/0011019 [math.PR]. URL: https://arxiv.org/abs/math/0011019.

[4] Noam Berger et al. "Asymptotic Behavior and Distributional Limits of Preferential Attachment Graphs". In: *The Annals of Probability* 42.1 (2014), pp. 1–40. DOI: 10.1214/12-AOP755.

[5] Vladimir I. Bogachev. *Measure Theory*. 2 vols. Berlin: Springer-Verlag, 2007. DOI: 10.1007/978-3-540-34514-5.

[6] Béla Bollobás. *Random Graphs*. 2nd ed. Cambridge University Press, 2001.

[7] D. J. Daley and D. Vere-Jones. *An Introduction to the Theory of Point Processes, Volume I*. 2nd ed. Probability and Its Applications. Springer, 2003.

[8] Donald Dawson and Luis Gorostiza. *Percolation in an ultrametric space*. 2013. DOI: 10.48550/arXiv.1006.4400. arXiv: 1006.4400 [math.PR]. URL: https://arxiv.org/abs/1006.4400.

[9] Reinhard Diestel. *Graph Theory*. 5th ed. Springer, 2017.

[10] David S. Dummit and Richard M. Foote. *Abstract Algebra*. 3rd ed. Wiley, 2004.

[11] Eric O. Endo and Daniel Valesin. "Local Limits of Spatial Gibbs Random Graphs". In: *ALEA, Latin American Journal of Probability and Mathematical Statistics* 17 (2020), pp. 51–63. DOI: 10.30757/ALEA.v17-03. URL: https://alea.math.cnrs.fr/articles/v17/17-03.pdf.

[12] Paul Erdős and Alfréd Rényi. "On Random Graphs I". In: *Publicationes Mathematicae* 6 (1959), pp. 290–297.

[13] Brian H. Fralix and Germán Riaño. "A new look at transient versions of Little's law, and M/G/1 preemptive last-come-first-served queues". In: *Journal of Applied Probability* 47.2 (2010), pp. 459–473. DOI: 10.1239/jap/1276784903. URL: https://www.cambridge.org/core/services/aop-cambridge-core/content/view/F56027639A3A8F20758AC74636D6090D/S0021900200006756a.pdf/a-new-look-at-transient-versions-of-littles-law-and-mg1-preemptive-last-come-first-served-queues.pdf.

[14] Geoffrey Grimmett. *Percolation*. 2nd ed. Berlin: Springer, 1999.

[15] Gunnar Hellmund, Michaela Prokešová, and Eva B. Vedel Jensen. "Lévy-based Cox point processes". In: *Advances in Applied Probability* 40.3 (2008), pp. 603–629. DOI: 10.1239/aap/1222868178. URL: https://www.cambridge.org/core/services/aop-cambridge-core/content/

view/3BD0C43A354F909ABF3F98B641A6690F/S0001867800002718a.pdf/levybased_cox_point_processes.pdf.

[16] Remco van der Hofstad. *Random Graphs and Complex Networks, Vol. 1*. Book draft / lecture notes. 2018. URL: http://www.win.tue.nl/~rhofstad/NotesRGCN.pdf.

[17] Remco van der Hofstad. *Random Graphs and Complex Networks, Vol. 2 (in preparation)*. Lecture notes. 2021. URL: http://www.win.tue.nl/~rhofstad/NotesRGCNII.pdf.

[18] Remco van der Hofstad, Pim van der Hoorn, and Neeladri Maitra. *Local limits of spatial inhomogeneous random graphs*. 2022. arXiv: 2107.08733 [math.PR].

[19] Olav Kallenberg. *Foundations of Modern Probability*. 2nd ed. Probability and Its Applications. Springer, 2002.

[20] Olav Kallenberg. *Random Measures, Theory and Applications*. Cham: Springer, 2017. DOI: 10.1007/978-3-319-41598-7.

[21] Alexander S. Kechris. *Classical Descriptive Set Theory*. Vol. 156. Graduate Texts in Mathematics. New York: Springer, 1995. DOI: 10.1007/978-1-4612-4190-4.

[22] Achim Klenke. *Probability Theory: A Comprehensive Course*. 3rd ed. Springer, 2020.

[23] Günter Last. *Modern random measures: Palm theory and related models*. Preprint, Fakultät für Mathematik, Universität Karlsruhe, 2008-13. 2008. DOI: 10.5445/IR/1000010948. URL: https://publikationen.bibliothek.kit.edu/1000010948/1023386.

[24] Günter Last. *Stationary random measures on homogeneous spaces*. Preprint, Fakultät für Mathematik, Universität Karlsruhe, 2008-16. Journal version: Journal of Theoretical Probability 23(2):478–497 (2010), doi:10.1007/s10959-009-0231-9. 2008. DOI: 10.5445/IR/1000012083. URL: https://publikationen.bibliothek.kit.edu/1000012083/1022543.

[25] Günter Last and Mathew D. Penrose. *Lectures on the Poisson Process*. Cambridge University Press, 2017.

[26] Günter Last and Hermann Thorisson. "Invariant transports of stationary random measures and mass-stationarity". In: *The Annals of Probability* 37.2 (2009), pp. 790–813. DOI: 10.1214/08-AOP420. URL: https://arxiv.org/pdf/0906.2062.pdf.

[27] Frédéric Lavancier, Jesper Møller, and Ege Rubak. "Determinantal Point Process Models and Statistical Inference". In: *Journal of the Royal Statistical Society: Series B (Statistical Methodology)* 77.4 (2015), pp. 853–877. DOI: 10.1111/rssb.12096.

[28] Jesper Møller and Rasmus Plenge Waagepetersen. *Statistical Inference and Simulation for Spatial Point Processes*. Boca Raton: Chapman & Hall/CRC, 2003.

[29] Sidney A. Morris. *Topology Without Tears*. 2015.

[30] Benjamin Nehring. "Point Processes in Statistical Mechanics: A Cluster Expansion Approach". PhD thesis. Universität Potsdam, 2012. URL: http://opus.kobv.de/ubp/volltexte/2012/6268/.

[31] Jerzy Neyman and Elizabeth L. Scott. "Statistical Approach to Problems of Cosmology". In: *Journal of the Royal Statistical Society. Series B (Methodological)* 20.1 (1958), pp. 1–43. DOI: 10.1111/j.2517-6161.1958.tb00272.x.

[32] Zdzisław Pawlas. "Central Limit Theorem for Random Measures Generated by Stationary Processes of Compact Sets". In: *Kybernetika* 39.6 (2003), pp. 719–729.

[33] Alex Sánchez Flores. "Convergence of Spatial Inhomogeneous Random Graphs as Point Processes". Bachelor's Thesis. Eindhoven University of Technology, Aug. 4, 2025.

[34] Cosma Rohilla Shalizi. *Advanced Probability II or Almost None of It*. Lecture notes. 2006. URL: https://www.stat.cmu.edu/~cshalizi/754/2006/notes/all.pdf.

[35] Anush Tserunyan. *Introduction to Descriptive Set Theory*. Lecture notes. 2024. URL: https://homepages.math.uic.edu/~sinapova/Anush%20DST%20lecture%20notes.pdf.

[36] Victor Veitch and Daniel M. Roy. *The Class of Random Graphs Arising from Exchangeable Random Measures*. 2015. arXiv: 1512.03099 [math.ST]. URL: https://arxiv.org/abs/1512.03099.

[37] Florian Voss, Catherine Gloaguen, and Volker Schmidt. "Random tessellations and Cox processes". In: *New Perspectives in Stochastic Geometry*. Ed. by W. S. Kendall and I. Molchanov. Oxford: Oxford University Press, 2009, pp. 12–35. URL: https://www.mathematik.uni-ulm.de/stochastik/personal/schmidt/publications/cox.

[38] Quan Zhou. *STAT 614: Unit 18, Convergence of Probability Measures*. Lecture notes. Definition 18.1 and Proposition 18.2. URL: https://zhouquan34.github.io/MathProb/STAT614_Unit18_QuanZhou.pdf.