Benjamini-Schramm Convergence of Random Rooted Graphs

Gabriel Spessato Agostini

MATH GU4053 - Topology (Fall 2021)

Abstract

In this paper I discuss the local convergence (or Benjamini-Schramm convergence) of random graphs. I start with a motivation of pursuing a theory of graph convergence in order to study large real-world networks and network models. Then I present basic definitions and results in graph theory which are fundamental for the paper. In section 3, I construct the metric space of rooted graphs and present the notion of local convergence, which is expanded to random rooted graphs in section 4 with a plethora of (somewhat surprising) graph limits. I conclude by discussing some network indicators that can be well-approximated or bounded under local convergence, justifying the pursuit of such limit theory.

Audience: I write to a student familiar with point-set topology, although a highly motivated student of real analysis will do just fine. Knowledge of convergence in metric spaces is essential. Formal treatement of the content on random graphs also requires knowledge of Probability Theory, including basics of random variable convergence and measurable metric spaces. I opt for covering that section with less rigour and more visual intuition, but additional references are suggested for the interested reader. No graph theory pre-requisite is assumed and the necessary background is provided.

Contents

1 Motivation

Much of contemporary information is structured in the form of networks. That is, a set of objects with respective qualities is interneconneted by particular relationships. In the internet, webpages link to one another; streets of a city connect their intersections; and people's acquaintances determine a network in society. Processes and games involving these relationships are illuminated by Graph Theory, the branch of math which studies networks—or graphs, how mathematicians prefer to call them. Gossiping and epidemic spreading (the latter being an unfortunately timely application) are dynamic processes taking place on social networks, for example. The simplicity yet pervasiveness of networks in the real-world justifies the detailed study of Graph Theory and its applications.

With important exceptions such as physically constrained street networks, most real-world networks are untractably large. As a result, often times scientists focus their attention on local phenomena: what are the properties of specific nodes, how clustered are small communities, is it possible to partition the nodes evenly, etc [1]. To understand some of these indicators, a clever idea is to view the large real-world networks as limiting objects of increasingly large finite graphs. This sets forward the necessity of developing a Graph Limit Theory–that is, a topological structure on some set of networks that allows for meaningful notions of distance and convergence.

Beyond size, most real-world networks are also sparse. That is, there are significantly less relationships than objects. A common example is an online social network with plenty of users, but with relatively few friendship connections. Some Graph Limit theories work well for dense graphs–such is the case of graphons [2]–but fail to represent sparsity.

Real-world phenomena must also be studied with certain appreciation of uncertainty. Privacy concerns often restrain detailed access to social network data via call data records or social media interactions, for example. The best we can often do is assign a probability to the existence of some social interaction. Our limit theory must be robust to probabilistic objects, to random graphs.

In this paper, I present the Benjamini-Schramm convergence of graphs. Through constructing a compact, complete metric space, this notion of convergence gives rise to a graph limit theory that works well for sparse and random graphs from the local point of view.

2 Graph Theory Preliminaries

This section collects results and definitions in Graph Theory which will be directly referred to in the pursuit of a graph limit theory. More detailed proofs and expanded results can be found in introductory texts such as [3] and [4].

2.1 Basic Graphs

I begin by introducing key concepts in graph theory needed to translate these real-world network relationships into Mathematical objects.

Definition 1. *A* graph $G = (V, E)$ *consists of its* vertex set $V(G)$ *and its* edge set $E(G)$*, where* $E(G)$ *is a subset of the set of two-element subsets of* $V(G)$*.*

This definition of a graph implies that there is at most one edge between any two vertices and that there are no self-edges i.e. edges starting and ending at the same vertex. Some graph theorists call the above object a *simple* graph for these reasons. Additionally, the graphs I presented are *undirected* because there is no distinction between the two vertices of an edge.

Example. Draw any collection of line segments in a piece of paper. You just drew your first graph! The vertices are all the endpoints of the segments, and the edges are the segments themselves. □

Note that the lines in the above example need not to intersect each other. In fact, a collection of entirely disjoint lines is a well defined graph, as is a collection of several points and no lines (that would be a graph with no edges). We will get back to these ideas later.

Definition 2. *Let $G = (V, E)$ be a graph. Then any graph $H = (V', E')$ such that $V' \subseteq V$ and $E' \subseteq E$ is a* subgraph *of G.*

Example. Take your drawing of lines—assuming you did not crumble the paper yet. Erase any number of segments. You just found a subgraph of your first graph. □

The idea of subgraphs raises the possibility of defining a maximal graph:

Example. A *complete graph on n vertices* contains an edge between any pair of vertices. That is, it is a graph $K_n = (V, E)$ such that $|V| = n$ and $E = \binom{V}{2}$. □

Thinking back to the idea of sparseness in real-world networks, we could say that complete graphs are the most possible dense networks. Any person in a complete social network is acquainted to everyone else—this is obviously never true in real life. But we can think of any graph as a subgraph of a complete graph, as if out of all possible connections only some of them actually exist. An example of complete graphs and their subgraphs is illustrated in figure 1.

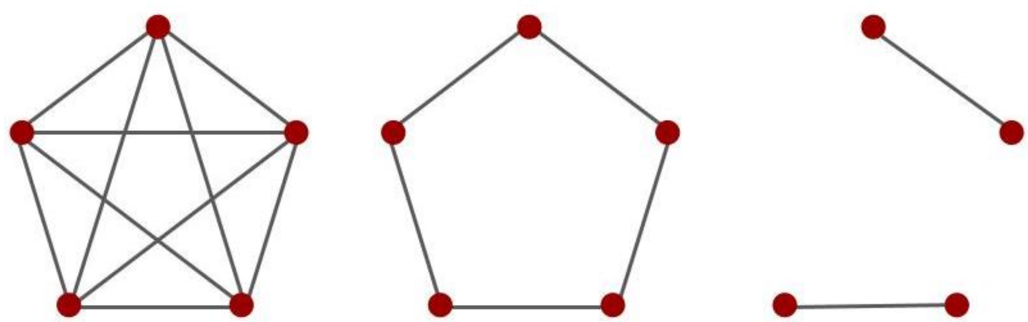

Figure 1: The complete graph on 5 vertices and two of its subgraphs, which have respectively 5 and 4 vertices each. Note that we can erase any edges or vertices to obtain a subgraph.

Using just the basic ideas of vertices and edges, we can define further measurements of adjacency:

Definition 3. *Let G be a graph. Let $x, y \in V(G)$ be vertices of G.*

1. We say that x and y are neighbors *if $\{x, y\} \in E(G)$. That is, x and y are neighbors if there is an edge between them.*

2. The degree $\deg(x)$ *of x is the number of neighbors of x. Equivalently, by the previous definition, it is the number of edges incident to x.*

*3. A graph in which all vertices have equal degree d is said to be d-*regular.

Example. The complete graph K_n is $(n-1)$-regular: any vertex is connected to all other vertices. □

2.2 Paths, Cycles, and Connectedness

This definition of graph is extremely rich. As we saw, any collection of points and lines in Euclidean space constitutes a graph. To help us produce some more examples and properly name them, we must build some vocabulary.

Definition 4. *Let G be a graph.*

1. A path of length k in G *is a sequence of vertices $x_0, x_1, \ldots, x_k \in V(G)$ such that x_i and x_{i+1} are neighborhs for all $i < k$. If $x_i \neq x_j$ for all $i \neq j$ we say that this is a* simple path.

2. A cycle of length k in G *is a path $x_0, x_1, \ldots, x_k$ such that $x_0 = x_k$ and $x_i \neq x_j$ otherwise.*

That is, a path is a sequence of adjacent vertices in a graph, as if you were tracing them without lifting the pencil. If there are no repeating vertices, then we have a simple path. And, if you finish exactly in the vertex you started, we have a cycle. The length of these objects is simply the count of edges they have.

Example. A *path graph of length* n and a *cycle graph of length* n are graphs P_n and C_n such that:

$$V(P_n) = \{v_0, v_1, \dots v_n\} \quad \text{and} \quad E(P_n) = \{\{v_k, v_{k+1}\} \mid k = 0, \dots n-1\}$$
$$V(C_n) = \{v_1, \dots v_n\} \quad \text{and} \quad E(C_n) = \{\{v_k, v_{k+1}\} \mid k = 1, \dots n-1\} \cup \{\{v_n, v_1\}\}$$

Note that P_n has n edges *but* $n+1$ *vertices*, whereas C_n has n edges and n vertices. The cycle graphs are always 2-regular, while in the path graphs there are two vertices of degree 1 and all other vertices have degree 2. □

Figure 2 illustrates the idea of paths and cycles.

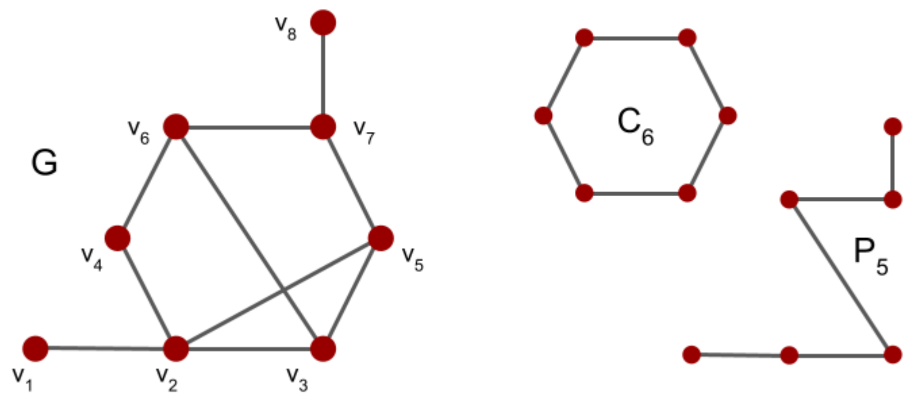

Figure 2: In the graph G, whose vertices are labelled $v_1 \dots v_8$, we can find multiple paths and cycles. For example, the sequence v_2, v_3, v_5 forms a cycle of length 3. The graphs C_6 and P_5 are, respectively, a cycle graph of length 6 and a path graph of length 5. Note that we can find such graphs also as subgraphs of G: the sequence $v_2, v_3, v_5, v_7, v_6, v_4$ forms a cycle of length 6 in G, and the sequence $v_1, v_2, v_3, v_6, v_7, v_8$ forms a path of length 5.

With the notion of paths, we can return to the unsettling idea that a collection of any non-intersecting lines in a plane is a graph. It turns out such graph is not minimal, in a sense.

Definition 5. *The graph* G *is* connected, *if for all* $x, y \in V(G)$ *there is a path from* x *to* y.

We can restrict our attention to these "nicer" connected graphs because all graphs have *connected components*: maximal connected subgraphs. To make this idea rigorous, consider the rightmost graph in 1. We can say that such graph is a *disjoint union* of two paths of length 1 (i.e. two edges, exactly by the definition) because there is no edge between these two graphs. The following lemma generalizes this notion:

Lemma 1. *Every graph* G *is the disjoint union of connected graphs, the* connected components *of* G.

Proof. Let $\sim$ be a binary relation on $V(G)$ such that $u \sim v \iff \exists$ a path from u to v. I show that $\sim$ is an equivalence relation on $V(G)$.

(i) u itself is a path of length 0 vacuously. So $u \sim u$ for all $u \in V(G)$.
(ii) if $u, x_1 ..., x_n, v$ is a path from u to v in G, then $v, x_n, ...x_1, u$ is a path from v to u,[1] So

[1] This works essentially because the graphs I defined are undirected! In the case of directed graphs there are additional notions of strong and weak connectedness.

$u \sim v \implies v \sim u$ for all $u, v \in V(G)$.
(iii) if $u, x_1..., x_n, v$ is a path from u to v in G and $v, y_1..., y_n, w$ is a path from v to w in G, then the concatenation $u, x_1..., x_n, v, y_1..., y_n, w$ is a path from u to w by the definition of path. So $u \sim v, v \sim w \implies u \sim w$ for all $u, v, w \in V(G)$.

Then $\sim$ induces a partition on $V(G)$. Let the equivalence classes with respect to $\sim$ be H_1, H_2.... We notice that $u, v \in H_i \iff \exists$ a path from u to v so each class is a connected graph itself. Moreover, there cannot be an edge between H_i and H_j for $i \neq j$: if there were, then there would be a path from a vertex of H_i to a vertex of H_j and they would have to be in the same equivalence class. Thus, the classes $H_1, H_2, ...$ partition not only the set of vertices of G but also the entire graph, and each of them is a connected graph. □

Connected graphs are then a basic object of study, since we can look at the components of any graph individually. The following is perhaps the most important type of graph we will consider:

Definition 6. *A* forest *is a graph with no cycles. A* tree *is a connected forest. The vertices of degree* 1 *in a tree are called* leaves.

Importantly, there is a unique path between any two vertices of a tree. The rigorous proof of this fact can be found in [3], but one can be convinced by noticing that the concatenation of two distinct paths with common endpoints defines a cycle. This also implies that, in a tree, any path is a path of shortest length.

Example. Figure 3 shows some trees. Note that the path graphs we have defined earlier are naturally trees as well. And the rightmost graph in figure 1 is a forest. □

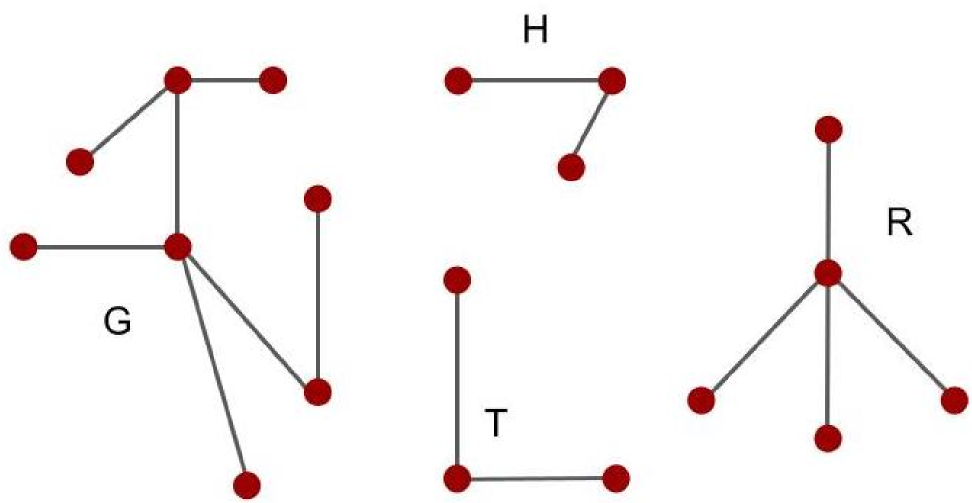

Figure 3: The graphs G, H, T, and S are all trees because they are connected and acyclical. In addition, the disjoint union of any collection of these graphs is a forest.

The idea of paths can lead to treating graphs as metric spaces. There is an inherent distance between two nodes determined by shortest path lengths. This distance turns out to be well-defined:

Definition 7. *For $x, y \in V(G)$ let the* distance $d(x, y)$ *be the minimal length of a path from x to y.*

We allow d to be infinite, which will happen in case x and y lie in two distinct components of G.

Lemma 2. *Let G be a connected graph. Then d is a metric on $V(G)$.*

Proof. First, notice that because G is connected there is a path between any two vertices $x, y \in V(G)$. Then there is also a path whose length is minimal, and no pair has infinite distance. So $d : V(G) \times V(G) \to \mathbb{R}$ is a well-defined function. We check d is a metric:

(i) $d(x, y) \geq 0$ for all $x, y \in V(G)$ by definition, since the length of a path is the (natural) number of edges it contains. There is a path with length zero from x to y, that is which does not go through any edge, if and only if $x = y$.

(ii) $d(x, y) = d(y, x)$ for all $x, y \in V(G)$, as a path from x to y can be inverted to yield a path from y to x with the same length and vice-versa.

(iii) $d(x, z) \leq d(x, y) + d(y, z)$ for all $x, y, z \in V(G)$, as concatenating a path from x to y and a path from y to z always yields a path from x to z whose length is equal to the sum of the lengths of the two paths. A minimal path from x to z has length at most the length of that specific path. □

Using this metric, we can generalize important notions of metric spaces to the context of a graph.

Definition 8. *A k-*ball *$B_k(x_0)$ in a connected graph G centered at some $x_0 \in V(G)$, is the set of vertices $v \in V(G)$ such that $d(x_0, v) \leq k$ and the set of edges between such vertices. Similarly, a k-*sphere *is the set of vertices $v \in V(G)$ such that $d(x_0, v) = k$ for a non-negative integer k.*[2]

The 1-sphere around a node is often referred to as its *neighborhood* and denoted $N(x_0)$.

Example. Figure 4 shows a graph and two balls of a certain graph. Note that $B_k(x_0) \subseteq B_{k+1}(x_0)$, a phenomenon that must be familiar to anyone who studied metric spaces. □

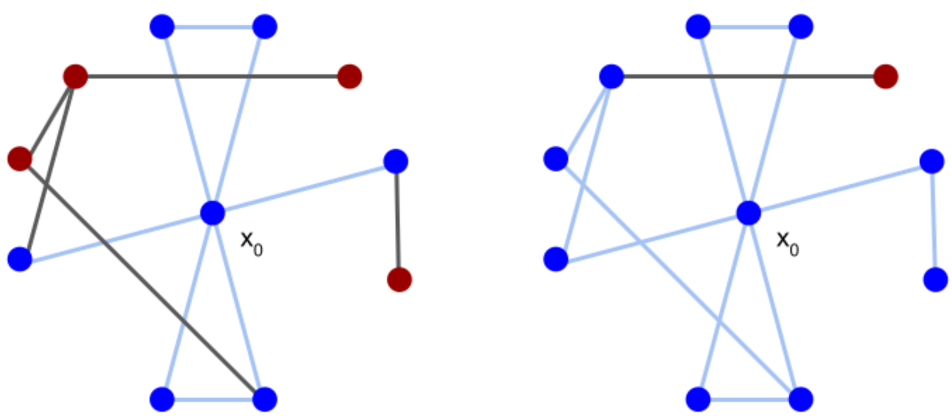

Figure 4: A graph with a distinct vertex x_0 and the 1-and 2-balls around that vertex in the left and right pictures respectively. Nodes and edges belonging to each ball are painted in blue.

2.3 Graph Isomorphisms

We now turn to defining a notion of equality in the set of graphs.

Definition 9. *Let G and H be two graphs. An* isomorphism *between G and H is a bijection $f : V(G) \to V(H)$ preserving adjacency relations. That is:*

$$\{u, v\} \in E(G) \iff \{f(u), f(v)\} \in E(H)$$

If there exists such a bijection, we say that G is isomorphic to H.

[2] It makes less sense to define the spheres as a graph–i.e. with vertices and edges–than as simply a set of vertices. But feel free to think of it as a subgraph of G, like the balls.

Graph isomorphisms are nothing but relabelings of the vertices that preserve neighbors. We can also think of them as two bijections that agree with each other, one between vertices and another between edges.

Example. In figure 3, the trees H and T are isomorphic. To build an isomorphism, the vertices of degree 2 must be mapped to each other. The other vertices can be mapped in any order. □

The following lemma is a logical step to justify the use of our definition of isomorphisms:

Lemma 3. *Graph isomorphism is an equivalence relation on the set of all graphs.*

Proof. Let $\sim$ be a binary relation on the set of all graphs such that $G \sim H$ if and only if G is isomorphic to H. Then:

(i) since the identity function on $V(G)$ is a bijection, $G \sim G$.
(ii) since isomorphim requires a bijection, we can use the inverse function f^{-1} to find an isomorphism from H to G.
(iii) if $G \sim H$ and $H \sim J$, then there exist adjacency-preserving bijections $f : G \to H$ and $g : H \to J$. The composite $g \circ f : G \to J$ is also a bijection and:

$$\{u, v\} \in E(G) \iff \{f(u), f(v)\} \in E(H) \iff \{g(f(u)), g(f(v))\} \in E(J)$$

So that $G \sim J$.

Then $\sim$ is reflexive, symmetric and transitive. Thus it is an equivalence relation. □

Example. Figure 5 shows an example of an isomorphism between graphs. Note that isomorphic graphs need to be "drawn" the same way in the paper, but that one could easily "redraw" them alike. Try doing that with the graphs in the figure to convince yourself. □

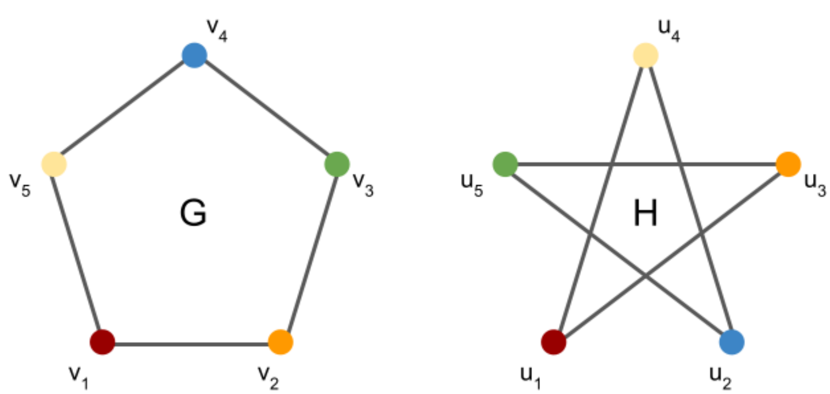

Figure 5: Isomorphic graphs G and H. A particular isomorphism f maps the vertices of G to the vertices of H such that vertices of equal color in the figure correspond. That is: $f(v_1) = u_1$, $f(v_2) = u_3$, $f(v_3) = u_5$, $f(v_4) = u_2$, and $f(v_5) = u_4$. Adjacency relationships are preserved.

In our limit theory, we will need a slightly stronger definition of isomorphism.

Definition 10. *Let G be a graph such that $v_* \in V(G)$.*

1. *The pair (G, v_*) is called a* rooted graph *and v_* is called the* root *of G.*
2. *The* radius *of (G, v_*) is the maximal distance from any vertex of G to v_*.*
3. *Two rooted graphs are isomorphic if there is an isomorphism identifying the roots.*

Example. In the example above, (G, v_4) and (H, u_2) are isomorphic as rooted graphs. Actually, G and H are isomorphic as rooted graphs independently of their roots: to see that, notice that we can build a satisfactory isomorphism by "shifting" the assignment of f until the roots agree. □

Example. In general, rooted isomorphisms are more restrictive. Figure 6 shows a case of isomorphism graphs which are not rooted isomorphic for a particular choice of root. □

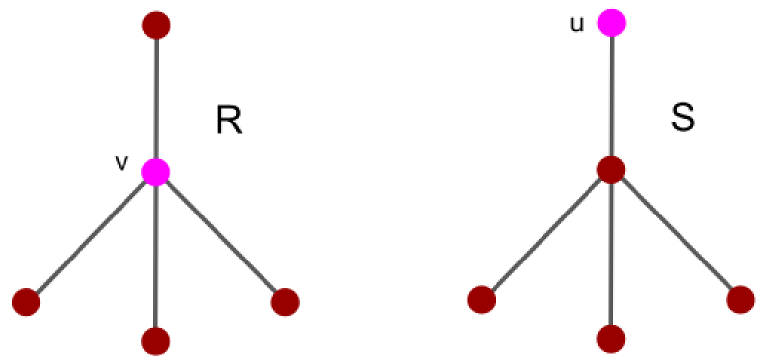

Figure 6: The trees R and S are isomorphic, simply match the vertices occupying the same position. However, the rooted trees (R, v) and (S, u) are not isomorphic. To see that, note that whereas v has degree 4 in R, u is a leaf of S and has degree 1. Any bijection from $V(R)$ to $V(S)$ mapping v to u will assign too many neighbors in R to u.

2.4 Infinite Graphs

The main objects we will be working with are infinite graphs. In particular, we restrict our attention to a subset of countable graphs:

Definition 11. *A graph is* locally finite *if all of its vertices have finite degree.*

Example. Any finite graph is locally finite. □

Example. Any d-regular tree must be an infinite graph as long as $d > 1$ and locally finite as long as $d < \infty$. To see that, let T be a d-regular tree and pick a vertex $v \in T$. By definition, v contains d neighbors, call them v_{11}, v_{12}, $\ldots$, v_{1d}. Then, each of these vertices contain d neighbors as well and, due to the absence of cycles, necessairly $d-1$ neighbors which have not been enumerated so far. We can keep enumerating vertices of T and we will always find $d-1$ new vertices so that T must have infinitely many of them. Verify this principle holds in the regular tress shown in figure 7. Hold on to the idea for this proof, we will make it rigorous in a while. □

The infinite path motivates another definition. Notice that there was no clear "starting point" in T_2. Let's addres this:

Definition 12. *Let $G = (V, E)$ be an infinite graph. A* ray *in G is an infinite sequence of distinct vertices $x_0, x_1, \dots \in V(G)$ such that $\{x_i, x_{i+1}\} \in E(G)$ for all $i \geq 0$.*

To sucessfully prove results regarding infinite graphs, we need some tools. While Zorn's lemma is very powerful and works in general settings, for locally finite graphs we can use a more visual and simple tool known as Konig's Infinity Lemma. I prove the lemma in what is often known as its tree variation:

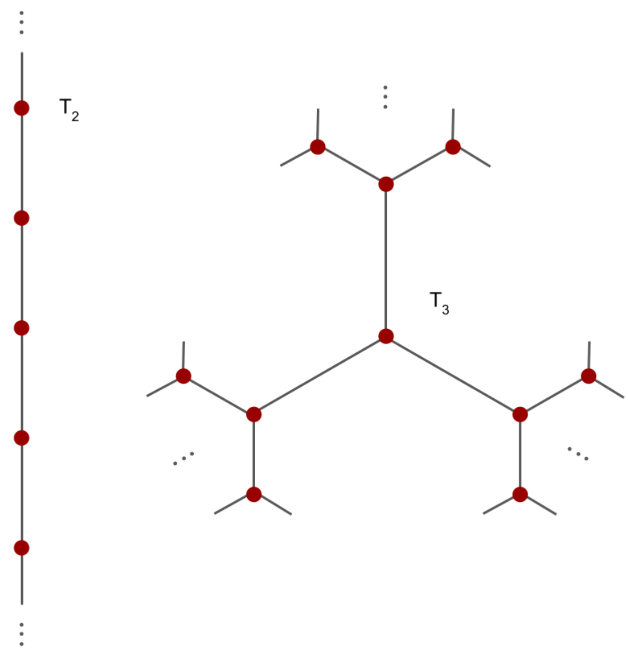

Figure 7: The 2-regular tree T_2 and the 3-regular tree T_3. Note that T_2 can be thought as an infinite path (and sometimes we will refer to it this way). These trees have no leaves. The number of vertices in the k-sphere of a tree grows exponentially.

Lemma 4. *Let T be an infinite tree rooted at x_0. If T is locally finite, then there exists a ray $x_0, x_1, \dots$ in T such that $d(x_0, x_n) = n$.*

Proof. To prove this lemma more smoothly we can use the definition of *descendants* in a rooted tree (T, x_0). The descendants of a vertex $v \in T$ are vertices of T whose (unique) simple path to the root passes through v. Now we use induction to prove that, for every natural number k, there exists a simple path $x_0, x_1, \dots x_k$ in T such that x_k has infinitely many descendants.

The base case amounts to picking x_0 alone. This is vacuously a simple path of length 0, and the descendants of the root are all the vertices–an infinite set by assumption.

Assume the result holds for $k \leq n$ and pick a simple path $x_0, \dots x_n$ where x_n has infinitely many descendants. Note that x_{n-1} is not a descendant of x_n since $x_0 \dots x_{n-1}$ is a simple path not passing through x_n. Therefore, x_n must have at least one neighbor which is not x_{n-1}, otherwise it would have no descendants.

Because the tree is locally finite, x_n has finitely many such neighbors. Enumerate them as $v_1, \dots v_d$. The descendants of these vertices (along with themselves) partition the set of descendants of x_n: for any descendant of x_n which is not its neighbor, a simple path to the origin must first pass through a neighbor since subsequent vertices in a path are adjacent. Thus we can partition the infinite set of descendants of x_n into sets of descendants of finitely many v_i.

By properties of set cardinality, at least one of such v_i must have infinitely many descendants. Pick this vertex and call it x_{n+1}. By the definition of a descendant, the simple path from x_{n+1} to the root passes through x_n and hence by uniqueness of simple paths in a tree it passes through $x_{n-1}, x_{n-2}, \dots x_1$. Therefore we found a simple path $x_0, \dots x_{n+1}$ where x_{n+1} has infinitely many descendants.

This completes the proof by induction. Note that, having found such a simple path for any natural number k, we found a ray by definition. $\square$

This result will be used in the construction of a metric space for rooted graphs. I will make direct use of the following corollary:

Corollary 1. *Let (G, v_0) and (H, u_0) be connected, countable, locally-finite rooted graphs. If the balls $B_r(v_0)$ and $B_r(u_0)$ are isomorphic for any $r \geq 0$, then G and H are isomorphic as rooted graphs.*

Rather than applying Konig's lemma directly to the graphs G or H themselves, to prove this result we will build a graph whose vertices are rooted isomorphisms between them. Then we show this new graph is a locally finite, infinite tree and hence we can use the infinity lemma to find a ray.

Proof. We build a graph W such that the vertices are rooted isomorphisms between all k-balls and the edges are restrictions between compatible isomorphisms:

$$V(W) = \bigcup_{k=0}^{\infty} V_k(W) = \bigcup_{k=0}^{\infty} \{f \mid f : B_k(v_0) \to B_k(u_0) \text{ is an isomorphism}\}$$

$$E(W) = \{\{f, g\} \mid \exists k : f \in V_k(W), g \in V_{k+1}(W), g|_{B_k(v_0)} = f\}$$

This graph is:

Connected. Any isomorphism can be restricted to the trivial isomorphism between the zero-balls $f_0 : \{v_0\} \to \{u_0\}$. Now consider W to be a rooted graph with root f_0.

Infinite. The balls $B_k(v_0)$ and $B_k(u_0)$ for all $k > 0$ are non-empty, because if they were the original graphs would not be infinite (since they are locally finite).

Locally finite. Each isomorphism $f \in V_k(W)$ has neighbors on $V_{k-1}(W)$ and $V_{k+1}(W)$ only by construction. Because the graphs G and H are locally finite, the k-balls are finite for any fixed $k \geq 0$. Between any two finite sets there is only a finite number of bijections—namely their cardinality factorial—so that there can only be finitely many isomorphisms between the k balls of G and H. This implies that both $V_{k-1}(W)$ and $V_{k+1}(W)$ are finite, and that f can have at most finitely many neighbors.

Acyclical. Any vertex $f_k \in V_k(W)$ has a unique neighbor in $V_{k-1}(W)$. That is because $B_{k-1}(x_0) \subseteq B_k(x_0)$ is a subgraph, therefore its assignments under f_k are fixed.

Therefore W is a locally-finite, infinite-tree rooted at f_0. By Konig's Lemma, there is an infinite ray in W. This infinite ray defines a rooted isomorphism between G and H: for any vertex of G we can find an assignment under an isomorphism (i.e. vertex) in W lying in the ray. $\square$

Konig's Infinity lemma can be used instead of Zorn's lemma to show that any countable graph contains a spanning tree, or that the chromatic number of a countable graph can be determined by the chromatic number of its finite subgraphs. These and other results can be found on chapter 8 of Diestel's text [3], and their proof follows a similar argument to the ball isomorphim proof above: building a tree based on restrictions of a function to finite graphs and applying Konig's lemma.

3 A Topological Space of Rooted Graphs

Let $\mathscr{G}$ denote the set of isomorphism classes of rooted, connected, locally-finite graphs. We start by putting a metric on $\mathscr{G}$. When not explicitly defined, assume that v_i is the root of a graph G_i.

Definition 13. *For $G_1, G_2 \in \mathscr{G}$ let*

$$d_{BS}(G_1, G_2) = \frac{1}{k+1}$$

where k is the maximal integer such that the rooted k-ball of G_1 is isomorphic to the rooted k-ball of G_2. If there is no such maximal integer then $d_{BS}(G_1, G_2) = 0$.

Lemma 5. *d_{BS} is a metric on $\mathscr{G}$.*

Proof. First, notice that $d_{\mathrm{BS}} : \mathscr{G} \times \mathscr{G} \to [0, \infty)$ is well-defined despite the choices of representatives G_1 and G_2 for the isomorphism classes. That is because isomorphic graphs have all their k-balls isomorphic, which is the only factor influencing d_{BS}. We now show identification, symmetry, and the triangular inequality:

(i) If $G_1, G_2 \in \mathscr{G}$ are isomorphic as rooted graphs, then any rooted k-ball of G_1 is isomorphic to a rooted k-ball of G_2. Therefore isomorphism implies that $d_{\mathrm{BS}}(G_1, G_2) = 0$. The converse follows from corollary 1, for if $d_{\mathrm{BS}}(G_1, G_2) = 0$ then all k-balls in G_1 and G_2 are isomorphic.

(ii) Due to the symmetry of graph isomorphisms, this distance is symetric.

(iii) We can prove the triangular inequality by proving the *strong triangular inequality*: any side of a triangle is no larger than the largest of the other two sides, hence any triangle is isosceles. This helps set up the Benjamini-Schramm space as a *ultrametric space*, an interesting result which will not be exploited in this paper.

Let $G_1, G_2, G_3 \in \mathscr{G}$. Let k_{ij} denote the maximal integer such that the rooted k_{ij}-ball of G_i is isomorphic to the rooted k_{ij}-ball of G_j (possibly infinity). That is:

$$\begin{aligned}
k_{ij} = \sup\{k \mid B_k^{G_i}(v_i) \cong B_k^{G_j}(v_j)\} \implies & \begin{cases} B_k^{G_1}(v_1) \cong B_k^{G_3}(v_3) & \forall k \leq \min\{k_{13}, k_{23}\} \leq k_{13} \\ B_k^{G_2}(v_2) \cong B_k^{G_3}(v_3) & \forall k \leq \min\{k_{13}, k_{23}\} \leq k_{23} \end{cases} \\
\implies & B_k^{G_1}(v_1) \cong B_k^{G_2}(v_2) \quad \forall k \leq \min\{k_{13}, k_{23}\} \\
\implies & k_{12} = \sup\{k \mid B_k^{G_1}(v_1) \cong B_k^{G_2}(v_2)\} \geq \min\{k_{13}, k_{23}\} \\
\implies & k_{12} + 1 \geq \min\{k_{13} + 1, k_{23} + 1\} \\
\implies & \frac{1}{k_{12}+1} \leq \max\left\{\frac{1}{k_{13}+1}, \frac{1}{k_{23}+1}\right\} \\
\implies & \frac{1}{k_{12}+1} \leq \frac{1}{k_{13}+1} + \frac{1}{k_{23}+1}
\end{aligned}$$

The second implication followed by transititviy of isomorphisms, and the third by the definition of a supremum. ∎

Endow $\mathscr{G}$ with the topology defined by d_{BS}. We call this metric the *Benjamini-Schramm space* after [5]. Considering the notion of convergence in metric spaces, we have established the Benjamini-Schramm convergence as the convergence of rooted balls:

$$\begin{aligned}
(G_n)_{n=1}^{\infty} \to_{\mathrm{BS}} G \quad & \iff \quad d_{\mathrm{BS}}(G_n, G) \to 0 \\
& \iff \quad \sup\{k \mid B_k^{G_n}(v_n) \cong B_k^{G}(v)\} \to \infty \\
& \iff \quad \forall k > 0\ \exists n_k \in \mathbb{N} : \ B_k^{G_n}(v_n) \cong B_k^{G}(v)\, \forall n > n_k
\end{aligned}$$

Now we take a look at three examples of the Benjamini-Schramm convergence:

Example. The sequence of paths P_n rooted at their left endpoint x_0 converges in the Benjamini-Schramm sense to an infinite one-sided ray rooted at its left endpoint, P. These graphs are illustrated in figure 8, and they satisfy the following relation:

$$d_{\mathrm{BS}}(P_n, P_m) = \frac{1}{\min(n, m) + 1} \quad \text{and} \quad d_{\mathrm{BS}}(P_n, P) = \frac{1}{n+1}$$

To see that this holds, observe that $P_n \subseteq P_m$ is a subgraph (up to isomorphism) whenever $n \leq m$. Therefore there is an isomorphism between the n-rooted ball of P_n (which is the whole graph) and the n-rooted ball of P_m. For P, think of it as a path when $n \to \infty$.

Therefore (P, x_0) is the Benjamini-Schramm limit of the $((P_n, x_0))_{n=1}^{\infty}$. Given any $\varepsilon > 0$ we can pick:

$$n > \frac{1}{\varepsilon} - 1 \implies d_{\mathrm{BS}}(P_n, P) < \varepsilon$$

$\square$

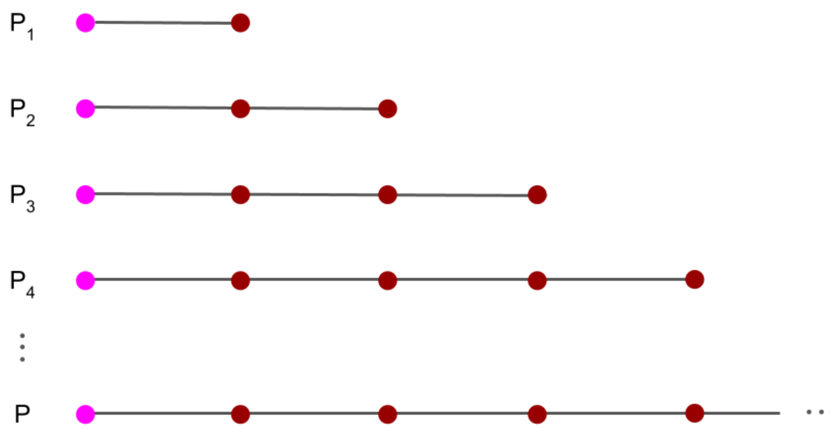

Figure 8: Paths (P_n, x_0) and their Benjamini-Schramm limit (P, x_0).

Example. The sequence of paths P_n rooted at their middle vertex $x_{\frac{1}{2}n}$ (or close to the middle if n is odd) converges in the Benjamini-Schramm sense to an infinite two-sided ray rooted anywhere, T_2, also known as the 2-regular tree from figure 7. These graphs are illustrated in figure 9, and they satisfy the following relation:

$$d_{\mathrm{BS}}(P_n, P_m) = \frac{1}{\left\lfloor \frac{\min(n,m)}{2} \right\rfloor + 1} \quad \text{and} \quad d_{\mathrm{BS}}(P_n, T_2) = \frac{1}{\left\lfloor \frac{n}{2} \right\rfloor + 1}$$

Again, this happens primarily because $P_n \subseteq P_m$ is a subgraph (up to isomorphism) whenever $n \leq m$. But in this case the n-ball in P_m may not be equal to P_n, since the root is in the middle of the graph (see the figure)! Note that the distance between any of these paths (except P_1) to P in figure 8 is 1, for their 1-balls are not isomorphic due to the degree of the root being different.

Therefore (T_2, x) is the Benjamini-Schramm limit of the $((P_n, x_{\frac{1}{2}}n))_{n=1}^{\infty}$. Given any $\varepsilon > 0$ we can pick:

$$n > \frac{2}{\varepsilon} - 2 \implies d_{\mathrm{BS}}(P_n, T_2) < \varepsilon$$

$\square$

Example. The sequence of cycle graphs C_n rooted at *any* vertex x_0 also converges in the Benjamini-Schramm sense to the infinite two-sided ray rooted at any vertex T_2. These graphs are illustrated in figure 10, and they satisfy the following relation:

$$d_{\mathrm{BS}}(C_n, C_m) = \frac{1}{\left\lfloor \frac{\min(m,n)}{2} - 1 \right\rfloor + 1} \quad \text{and} \quad d_{\mathrm{BS}}(C_n, T_2) = \frac{1}{\left\lfloor \frac{n}{2} - 1 \right\rfloor + 1}$$

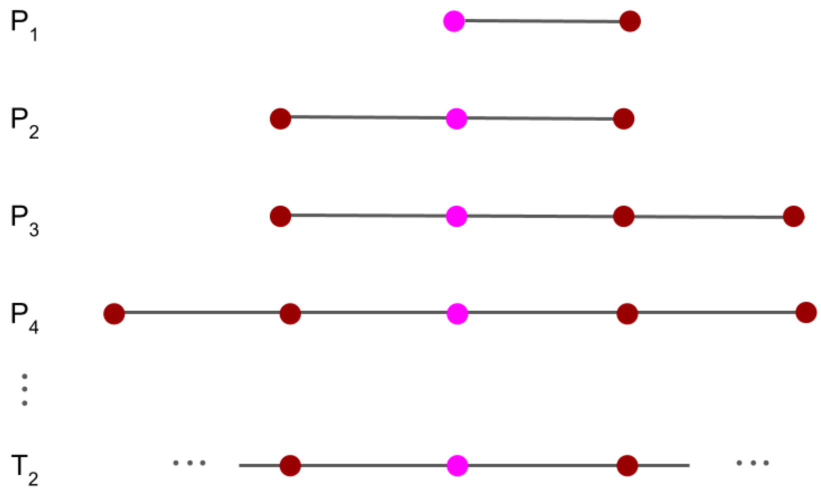

Figure 9: Paths $(P_n, x_{\frac{n}{2}})$ and their Benjamini-Schramm limit T_2. Note how the 1-ball in P_2 contains both neighbors of the root, hence is not isomorphic to the 1-ball in P_1. That is why, in the BS-distance of these graphs, we must take the floor of half the path length.

To see that this holds, observe that until we have transversed half the cycle (on each side of the root) we will not capture the edge "closing" the cycle. Then our balls will be isomorphic to a path rooted in its middle vertex such as the ones shown in figure 9 and we reduce to the previous example.

Therefore (T_2, v) is the Benjamini-Schramm limit of the $((C_n, v))_{n=3}^{\infty}$. Given any $\varepsilon > 0$ we can pick:

$$n > \frac{2}{\varepsilon} + 2 \implies d_{\mathrm{BS}}(C_n, T_2) < \frac{1}{\left\lfloor \frac{1}{\varepsilon} \right\rfloor + 1} \leq \frac{1}{\frac{1}{\varepsilon} - 1 + 1} = \varepsilon$$

□

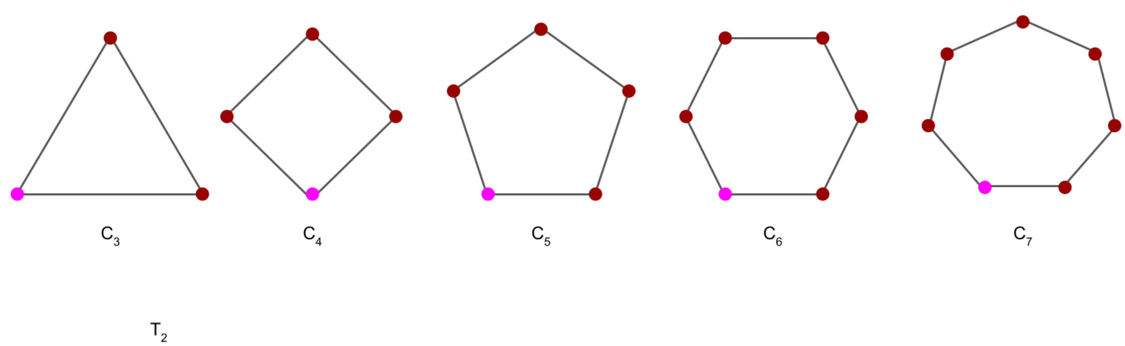

Figure 10: A few rooted cycles and their Benjamini-Scramm limit T_2. The reasoning for taking the floor is very similar to the one presented in 9. The correction to one unit less happens because in the case of an odd cycle the ball captures the entire graph before reaching half of its length. The 2-ball of C_5 is the entire pentagon.

There are three important takeways from these examples.

1. Benjamini-Schramm convergence is *local*. The cycles are clearly not akin to an infinite path (which is a tree!) globally, but from the point of view of the root they all look like a path. Think of the Earth (or, if you can, of a manifold) and how it feels flat from any point in it.

2. The roots matter. Rooting the sequence of paths at different vertices (with, importantly, different degrees) yields a different limit.
3. The roots are sometimes irrelavant. In the case of cycles, for all $v \in C_n$ the rooted graphs (C_n, v) are isomorphic. This isomorphism is easy to visualize: simply rotate the cycles.

To understand more how these graph limits work, we have to prove a few properties about the Benjamini-Schramm space.

3.1 Benjamini-Schramm is a Polish Space

A first result regarding this space is that infinite graphs can be well-approximated by sequencs of finite rooted graphs. This is analogous to how real numbers are limits of sequences of rational numbers, and is a consequence of separability:

Theorem 1. *$\mathscr{G}$ is a separable metric space, where the set of finite rooted graphs modulo isomorphism forms a dense subset.*

Proof. Let $\mathscr{F} \subseteq \mathscr{G}$ denote the set of equivalence classes of finite rooted graphs. For every natural number k there are finitely many graphs on k vertices (finitely many subgraphs of the complete graph) and k possible roots. As a countable union of finite sets, $\mathscr{F}$ is countable. We show it is dense.

Let $(G, v) \in \mathscr{G}$ be an arbitrary rooted graph. For all $k \in \mathbb{N}$, the k-ball $B_k^G(v)$ is a finite rooted graph hence an element of $\mathscr{F}$ which satisfies:

$$d_{BS}\left(B_k^G(v), G\right) \leq \frac{1}{k+1}$$

since the ball is a subgraph of G. Fix $\varepsilon > 0$. Then we can pick any k above $\varepsilon^{-1} - 1$ and:

$$d_{BS}\left(B_k^G(v), G\right) < \varepsilon$$

Which shows that $\mathscr{F}$ is dense in $\mathscr{G}$, since we picked $G \in \mathscr{G}$ arbitrarily and found a point in $\mathscr{F}$ arbitrarily close to it. $\square$

Therefore, instead of looking at an infinite graph (or, in the real-world, an untractably large graph), we can look at a sequence of finite graphs converging to it. But, as in any metric space, some sequences do not converge. Figure 11 shows an example of such sequence.

Fortunately, "nice" sequences converge. Recall the following definition from analysis:

Definition 14. *Let (X, d) be a metric space. A sequence $(x_n) \subseteq X$ is called a* Cauchy Sequence *if for all $\varepsilon > 0$ there exists $N = N(\varepsilon)$ such that if $n, m > N$ then $d(x_n, x_m) < \varepsilon$. The space X is said to be* complete *if all Cauchy Sequences converge.*

The divergent sequence in figure 11 was not Cauchy (pick $\varepsilon < 1$). In the Benjamini-Schramm space, Cauchy sequences converge. My proof follows the exposition in appendix A of [6].

Theorem 2. *$\mathscr{G}$ is a complete metric space.*

Proof. Let $((G_n, v_n)) \subseteq \mathscr{G}$ be a Cauchy sequence. First, let's see what the Cauchy condition means for the balls. Let $\varepsilon > 0$ be given and pick $N > 0$ as per the previous definition. Then:

$$\begin{aligned} d_{\mathrm{BS}}(G_n, G_m) < \varepsilon \ \forall n, m > N &\implies d_{\mathrm{BS}}(G_n, G_m) < \frac{1}{\left(\frac{1}{\varepsilon} - 1\right) + 1} \ \forall n, m > N \\ &\implies \sup\{k \mid B_k^{G_n}(v_n) \cong B_k^{G_m}(v_m)\} > \frac{1}{\varepsilon} - 1 \ \forall n, m > N \end{aligned}$$

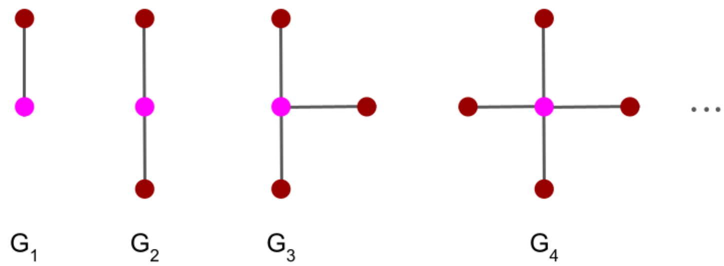

Figure 11: A sequence $(G_n) \subseteq \mathscr{G}$ rooted at the pink vertex, x_0. Algebraically, $V(G_n) = \{x_0, \dots x_n\}$ and $E(G_n) = \{\{x_0, x_i\}\}$. Note that the 1-rooted balls in any two of these graphs cannot be isomorphic: the root has degree n, different in all G_n, therefore no bijection between the 1-balls will preserve adjacency. This means that $k = 0$ and that the distance between any two elements of this sequence is 1. Thus it cannot converge to a point in $\mathscr{G}$, which makes sense: the logical limiting graph would contain infinitely many vertices, all adjacent to the root—but then this graph would not be locally finite!

$$\implies B_k^{G_n}(v_n) \cong B_k^{G_m}(v_m)\ \forall k \leq \frac{1}{\varepsilon} - 1\ \forall n, m > N$$

Equivalently, if we assume that k is given, we pick $\varepsilon < \frac{1}{k+1}$ and the associated N to find that

$$B_k^{G_n}(v_n) \cong B_k^{G_m}(v_m)\ \forall n, m > N$$

But this means that we can find a $n_k := N > 0$ such that

$$B_k^{G_n}(v_n) \cong B_k^{G_{n_k}}(v_{n_k})\ \forall n \geq n_k \qquad (*)$$

Using this reformulation of the Cauchy property, we will construct a sequence of graphs all within a given distance from an element of $\mathscr{G}$. This element will become the limiting object.

For every $k = 1, 2, \dots$, find one such n_k as given by (*). This can be done such that the sequence (n_k) is monotonically increasing: any $n > n_k$ also satisfies the given property. We construct a sequence of graphs using the k-balls:

$$\left(B_k^{G_{n_k}}(v_{n_k}), v_{n_k}\right) = (H_k, v_k) \subseteq \mathscr{G}$$

Note that, by (*) and the monotonicity of n_k, we have a sequence where:

$$B_r^{H_s}(v_s) \cong B_r^{H_r}(v_r) \quad \forall\, r \leq s$$

Therefore these graphs are compatible as rooted graphs: the r-balls freeze. We can consider that they have the same set of vertices (one could define a proper isomorphism). Let $(G, v) \in \mathscr{G}$ be another graph whose vertices are the union of all vertex sets $\cup V(H_k)$ and the edges are the union of all edge sets. Clearly G is locally-finite and connected since so are all the H_k. And, by the definition of G and (*), we have:

$$B_k^G(v) \cong B_k^{H_k}(v_k) = B_k^{G_{n_k}}(v_{n_k}) \cong B_k^{G_n}(v_n) \quad \forall\, n \geq n_k$$

This, in terms of the Benjamini-Schramm distance, tells us that:

$$d_{\mathrm{BS}}(G, G_n) \leq \frac{1}{k+1} \quad \forall\, n \geq n_k$$

This completes the proof because k was an arbitrary natural number. Thus, given any $\varepsilon > 0$:

$$k > \frac{1}{\varepsilon} - 1 \implies d_{\mathrm{BS}}(G, G_n) < \varepsilon$$

That is, $(G_n)_{n=1}^{\infty} \to_{\mathrm{BS}} G$. $\square$

Therefore, the Benjamini-Schramm topological space is a complete and separable metric space. We can say it is a *Polish space.*

3.2 Compactness in the Benjamini-Schramm Space

Take a look again at the sequence given in figure 11. As explained in the figure, the distance between any two elements of that sequence is 1. So, picking a few of those graphs, even if we choose them wisely, will not generate a convergent sequence. Not only the sequence diverges, but it also contains no convergent subsequence. If we recall the definition of (sequential) compactness in the setting of metric spaces:

Definition 15. *Let (X, d) be a metric space. If every sequence $(x_n) \subseteq X$ contains a convergent subsequence $(x_{n_k}) \subseteq (x_n)$, we say that X is* compact.

Therefore we showed by a counter-example that the Benjamini-Schramm space $\mathscr{G}$ is not compact. Compactness is a very appealing property, for continuous functions in compact spaces are well-behaved. It turns out we can easily characterize compact subspaces of $\mathscr{G}$. We begin by defining an important subspace.

Definition 16. *Let $D > 0$ be an integer. The subspace of $\mathscr{G}$ with locally-finite, connected rooted graphs (modulo isomorphism) of all degrees at most D is $\mathscr{G}_D \subseteq \mathscr{G}$.*

This subspace is compact. Before we show that, we can show two other important properties.

Theorem 3. *The metric space $\mathscr{G}_D$ is closed in $\mathscr{G}$.*

Proof. We show that $\mathscr{G}_D$ by showing it contains all of its limit points. Let $(G_n, v_n) \subseteq \mathscr{G}_D$ be a convergent sequence whose limit is $(G, v) \in \mathscr{G}$. We can prove that the degrees in G are all at most D by contradiction.

Assume there is a vertex $x \in G$ such that $\deg(x) > D$. Let $r = d^G(x, v)$ be the distance *in the graph G* between x and the root. Then $B^G_{r+1}(v)$ contains x and all of its neighbors, whose distance to the root is either $r - 1$ or $r + 1$. But this implies that, for any $(H, u) \in \mathscr{G}_D$, there cannot be a rooted isomorphism between $B^H_{r+1}(u)$ and $B^G_{r+1}(v)$: while in the former all the vertices have degree at most D, the latter contains the vertex x whose degree is more than D. Therefore, by definition of the Benjamini-Schramm distance:

$$B^H_{r+1}(u) \not\cong B^G_{r+1}(v) \implies d_{\text{BS}}(G, H) > \frac{1}{r+1+1} = \frac{1}{r+2} \quad \forall (H, u) \in \mathscr{G}_D$$

If we pick $\varepsilon < \frac{1}{r+2}$, there is no element of the sequence (G_n, v_n) which is within ε from G. Then the sequence cannot converge to G, which is a contradiction. So no vertex in G can have degree above D which implies that $G \in \mathscr{G}_D$.

The proof is complete since G was an arbitrary limit point of $\mathscr{G}_D$. □

Now recall the following definition from analysis:

Definition 17. *Let (X, d) be a metric space. We say that X is* totally bounded *if, for any $r > 0$,*

$$\exists x_1, x_2, \ldots x_{N_r} \in X : \exists i \in \{1, \ldots N_r\} : d(x, x_i) < r \quad \forall x \in X$$

I remark that usually the definition of total boundedness is phrased in terms of open balls. I avoided this since we are making use of balls in the graphs a lot.

Theorem 4. *The metric space $\mathscr{G}_D$ is totally bounded.*

Proof. Let $r > 0$ be given. We will use the definition of the radius of a rooted graph. Pick:

$$\rho > \left\lceil \frac{1}{r} \right\rceil$$

There are only finitely many possible graphs in $\mathscr{G}_D$ whose radius is at most ρ. This happens because such graph has at most finitely many, namely D^ρ, vertices. Call these graphs $G_1, \dots G_N$. We will show that any graph in $\mathscr{G}_D$ is within r from at least one of these graphs.

Let $(H, u) \in \mathscr{G}_D$ be an arbitrary rooted graph. The ρ-rooted ball of H, $B_\rho^H(u)$, is a a graph in $\mathscr{G}_D$ whose radius is at most ρ. Therefore it is isomorphic to one of the $G_1, \dots G_N$, say G_k. By definition of the Benjamini-Schramm distance, we then have:

$$B_\rho^H(u) \cong G_k \implies d_{\mathrm{BS}}(H, G_k) \leq \frac{1}{\rho + 1} < \frac{1}{\left\lceil \frac{1}{r} \right\rceil + 1} \leq \frac{1}{\frac{1}{r} + 1} \leq r$$

This completes the proof since r was arbitrary. □

We can now put these two properties together. The following lemma, from analysis, explains how we will do that:

Lemma 6. *A complete and totally bounded metric space is compact.*

In addition, recall from analysis that a closed subset of a complete space is also complete. This allows us to show:

Theorem 5. *The metric space $\mathscr{G}_D$ is compact.*

Proof. This is a consequence of theorems 3 and 4. Because $\mathscr{G}$ is complete and $\mathscr{G}_D \subseteq \mathscr{G}$ is closed, it is also complete. Then, following lemma 6, it is compact. □

Much of the local limits of interest happen to be confined to the subspace $\mathscr{G}_D$. The convergent sequences of paths and cycles shown in this section were all within $\mathscr{G}_2$ for example. The divergent sequence in figure 11 was not in any such subspace: for any choice of D, the root degree in graph G_D would exceed the bound. Appealing to compactness will allow us to easily generalize the Benjamini-Schramm convergence to the case of random graphs, as I will discuss next.

4 A Topological Space of Random Rooted Graphs

Often, real-world networks do not have a distinct vertex. To make sense of the idea of a rooted graph, we can think of choosing the root uniformly at random among a graph's vertices. This inspires the following definition:

Definition 18. *For a fixed integer $D > 0$ and the Borel sigma-algebra, a $\mathscr{G}_D$-valued random variable $\boldsymbol{G}$ is called a* random rooted graph*.*

Given a fixed finite graph, we can construct a random rooted graph naively by choosing its root uniformly. Note that sometimes we choose the root so that the resulting graphs are isomorphic. This yields a skewed probability distribution.

Example. Let P_n represent the path graph of length n with vertices $x_0, \dots x_n$. We associate to it a random rooted graph $\boldsymbol{P_n}$. The root will fall on vertex x_i with probability $\frac{1}{n}$ regardless of i. However, if we want to consider how this graph will look like up to isomorphism (hence, in $\mathscr{G}_D$), we can consider it to be rooted in half of its vertices:

(i) If n is odd, $\boldsymbol{P_n}$ is rooted at x_i with probability $\frac{2}{n}$ for any $i < \frac{n+1}{2}$.

(ii) If n is even, $\boldsymbol{P_n}$ is rooted at $x_{\frac{n}{2}}$ with probability $\frac{1}{n}$ and at another x_i, $i < \frac{n}{2}$ with probability $\frac{2}{n}$.

This happens because of symmetry: as rooted graphs, the path rooted at x_0 is isomorphic to the path rooted at x_n and so on. □

Example. Let C_n represent the cycle graph of length n. As stated in a takeway of the previous section, rooting a cycle at any of its vertices yields an isomorphic rooted graph. Therefore we can say that, with C_n, we associate a random rooted graph $\boldsymbol{C_n}$ whose root is a vertex v with probability 1. □

A random rooted graph, then, corresponds to a Borel probability distribution on $\mathscr{G}_D$. For a more detailed and rigorous probabilistic explanation, the reader can look at [7] and [8]. The Benjamini-Schramm convergence can be extended to random rooted graphs as a weak convergence (or convergence in probability) of random variables. That is:

$$(\boldsymbol{G_n})_{n=1}^{\infty} \to_{BS} \boldsymbol{G} \quad \iff \quad \forall (H, o) \in \mathscr{G}_D \ \Pr\left[B_k^{G_n}(v_n) \cong B_k^H(o)\right] \to \Pr\left[B_k^G(v) \cong B_k^H(o)\right]$$

Essentially, we are comparing the k-ball in each of the G_n to a k-ball in a fixed rooted graph H to make sense of the random variables. The transitiviy of graph isomorphism should help the reader understand why this notion of convergence implies that, in probability, the limit of $\boldsymbol{G_n}$ will be isomorphic to $\boldsymbol{G}$. More intuitively, we can think of how the balls (for a fixed radius) in G_n will look like if we choose a root at random. What kind of vertices will lie inside the ball? What graph are these balls converging to as $n \to \infty$? These questions illustrate once again how the Benjamini-Schramm focuses on the *local* structure of the graph, since convergence is determined by the balls.

This notion of convergence works well because our metric space $(\mathscr{G}_D, d_{BS})$ is compact and Polish (separable and complete). Prokhorov's theorem, an important result in Probability Theory, guarantees that the space of Borel-probability distributions on $(\mathscr{G}_D, d_{BS})$ is also compact. More details about Prokhorov's theorem can be found in [9].

4.1 Examples of Random Graph Convergence

Now I give several examples of the Benjamini-Schramm convergence of random rooted graphs:

Example. Convergence of cycles is once again a trivial case. The limiting object of the sequence $\boldsymbol{C_n}$ will be the two-dimensional tree $\boldsymbol{T_2}$ rooted, with probability 1, at any of its vertices. That is because, for fixed r, balls around a vertex of the cycle mostly look like balls in a two-sided path—considering, of course, the tail of the sequence of cycles. □

Example. Random rooted graphs solve the conflict between the two path limits. Let $\boldsymbol{P_n}$ be the random rooted path of length n, whose distribution was precisely defined on a previous example. I claim that the limit of $(\boldsymbol{P_n})$ is also $\boldsymbol{T_2}$ rooted, with probability 1, at any of its vertices.

This happens because the left endpoints become negligible as the paths grow. While the probability that the path P_n will be rooted at a vertex of degree 1 (i.e. endpoint) is $\frac{2}{n}$, the probability that it will be rooted at a vertex of degree 2 is $\frac{n-2}{n}$. As $n \to \infty$, the former approaches zero and the latter 1.

Thinking of balls, this means that the balls of fixed radius at a random vertex of P_n mostly extend nicely in both directions. That is, there is a very low probability a ball of fixed radius will capute the end point of the path as $n \to \infty$. □

In these two examples, the limit object had a degenerate probability distribution: with probability 1, the root was fixed. This is not usually the case: often the Benjamini-Schramm limit of random rooted graphs is random itself. These proofs follow the visual explanation in [10] and [11].

Example. Let F_n denote a *comb graph with handle* on $3n$ vertices. That is:

$$V(F_n) = \{b_1, \dots b_n\} \cup \{t_1, \dots t_n\} \cup \{h_1, \dots h_n\}$$
$$E(F_n) = \{\{b_i, b_{i+1}\} \mid i < n\} \cup \{b_n, h_1\} \cup \{\{h_i, h_{i+1}\} \mid i < n\} \cup \{\{t_i, b_i\} \mid i \le n\}$$

The vertices $t_1 \dots t_n$ are the teeth of the comb, the b vertices are the base of the comb, and the h vertices are the handle. It will be useful to consider the *comb graph without handle* M_n as the subgraph of F_n on the $2n$ vertices composing the base and teeth of the comb. This sequence of graphs is shown at figure 12.

As a sequence of random rooted graphs, $(\boldsymbol{F_n})$ converges to a random graph $\boldsymbol{F}$ which is isomorphic to, with equal $\frac{1}{3}$ probability:

(i) The infinite comb without handle M rooted at one of the base vertices.

(ii) The infinite comb without handle M rooted at one of the teeth vertices.

(iii) The infinite rooted 2-tree T_2.

This happens because $\frac{1}{3}$ of the vertices of F_n lies in each of the three regions of the comb. Taking the root uniformly over all vertices implies the root will be a tooth, a base vertex, or a handle vertex with equal probability. There are a few boundary cases: the points where handle meets the base, the endpoint of the handle, etc. As $n \to \infty$, there is zero probability that a ball in F_n will capture these finitely many boundary points. Locally, the graph looks like an infinite handleless comb or an infinite path—and in the case of the comb it matters where the root falls. $\square$

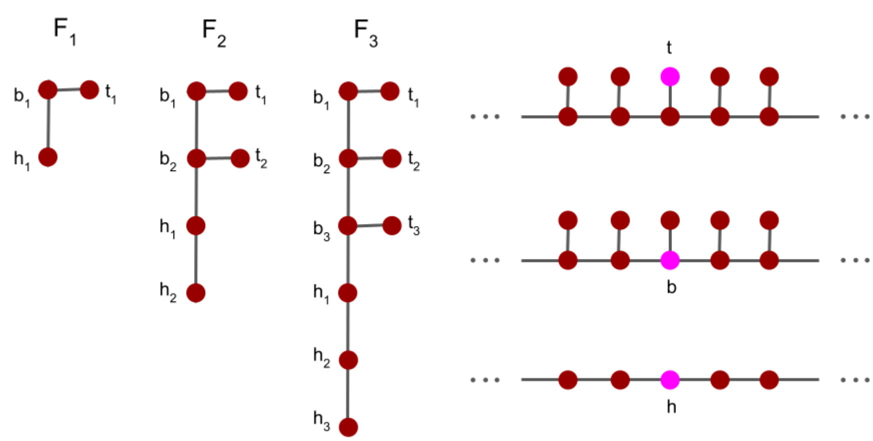

Figure 12: The (unrooted) comb graphs with handle F_1, F_2, and F_3 on the left and the realizations of their Benjamini-Schramm limit (as random rooted graphs) on the right. The graphs F_n have equal probability to be rooted on the teeth, base, or handle. Each of these scenarios corresponds to a value in the image of their limiting object.

There are many other interesting limits. Naturally extending the path to a k-dimensional lattice, for example, will yield that finite grids converge to the infinite grid. And readers who are more familiar with random graphs and stochastic processes will be pleased to know that Erdos-Renyi random graphs converge in the Benjamini-Schramm sense to the Galton-Watson tree with Poisson offspring distribution, a proof of which can be found in section 1.3 of [8]. I finish this section by illustrating one more interesting limit, whose limiting graph is perhaps unfamiliar to the reader.

Example. The sequence of truncated regular binary trees converges to the infinite canopy tree. Say that G_n is the binary tree which has $2^n - 1$ vertices. The first few of these trees are illustrated in figure 13.

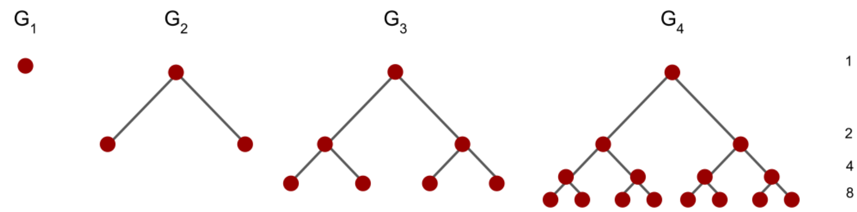

Figure 13: The sequence of truncated binary trees. Each tree G_n has $2^n - 1$ vertices distributed along n heights (consider the leaves are in height n i.e. we are counting down). And each height k of the tree has 2^{k-1} vertices. So G_n has 2^{n-1} leaves, which means that over half of the vertices in a truncated binary tree have degree 1.

The key to figuring out what is the limiting graph resides on examining how the tree looks, locally, from the perspective of its leaves, of the neighbors of its leaves, of the neighbors of the neighbors of its leaves, and so on. Rooting the tree at any vertex of the same height will yield an isomorphic rooted graph: think of it as if you are twisting the branches of the tree. We must root the truncated trees at every height. An equivalent visualization of the truncated tree G_4 is shown in figure 14.

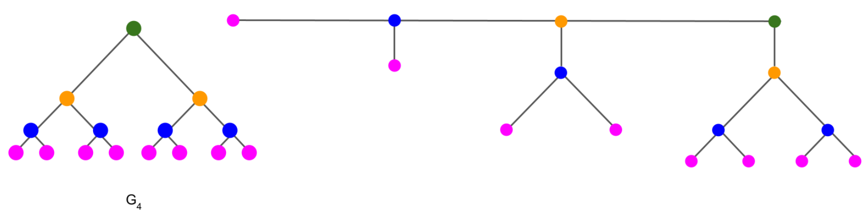

Figure 14: The truncated binary tree of height 4 in its common form (left) and on an isomorphic form (right). The colors map the vertices between the two graphs, such that vertices of the same color can be permuted without changing the isomorphism class. Basically, G_4 can be seen as the union of a 3 path and the truncated trees G_1, G_2 and G_3—each hanged from a vertex of the path. Rooting G_4 at a leaf is equivalent to rooting the figure on the right on the first vertex of the path, and so on.

It turns out that this construction generalizes well and the infinite graph it determines is called the *canopy tree*. The canopy tree is shown in figure 15.

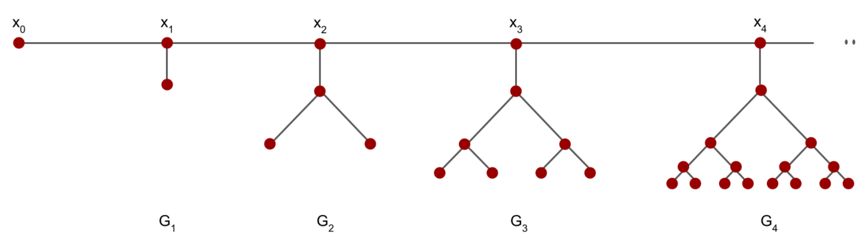

Figure 15: The canopy tree T_C. This infinite graph is composed of an infinite one-sided ray $x_0 x_1, \ldots$ and all truncated binary trees $G_1, G_2, \ldots$. The tree G_n hangs from the vertex x_n. Consider, if you wish, that G_0 is the empty graph.

Let's look at the balls rooted around a leaf. The truncated trees G_n are a subgraph of the canopy tree T_C. In their path-hanging form presented in figure 14, we can capture the whole height of G_n in a ball of T_C without noticing the two graphs are not isomorphic. But we notice it as soon as we increase the ball's radius one more unit, because the truncated path-hanging tree has a right endpoint. In the specific example of G_4 from figure 14, the 3-ball around a leaf will capture x_3 and that is ok, but the 4-ball will capture x_4 which has no correspondence in G_4 since the green vertex has no second neighborh in the path. The height of the tree G_n (in number of edges) is $n-1$. Generalizing this notion:

$$d_{\mathrm{BS}}((G_n, \text{leaf}), (T_C, x_0)) = \frac{1}{n-1+1} = \frac{1}{n}$$

In the deterministic Benjamini-Schramm sense, the sequence G_n when all graphs are rooted on a leaf converges to (T_C, x_0). Similarly, balls rooted around a vertex which is a neighbor of a leaf (the blue vertices in figure 14) are isomorphic for one less unit—since the blue vertices are one unit closer to the green vertex. This would mean:

$$d_{\mathrm{BS}}((G_n, \text{leaf neighbor}), (T_C, x_1)) = \frac{1}{n-2+1} = \frac{1}{n-1}$$

Again, as $n \to \infty$, we have found the deterministic Benjamini-Schramm limit.

Different than in the comb example, it seems like the limit of $\boldsymbol{G_n}$ will take infinitely many values with positive probability. They are all isomorphic to the canopy tree if we ignore the rooting. Their isomorphic class as rooted graphs, however, depends on the height of the original root, which was decided uniformly at random over the vertices of G_n. And, as given in figure 13, we know how to compute the limiting proportion of vertices at each height:

$$\begin{aligned}
\#(\text{leaves}) &= \frac{2^{n-1}}{2^n-1} \to \frac{1}{2} \\
\#(\text{neighbors of leaves}) &= \frac{2^{n-2}}{2^n-1} \to \frac{1}{4} \\
\#(\text{neighbors of neighbors of leaves}) &= \frac{2^{n-3}}{2^n-1} \to \frac{1}{8} \\
&\vdots
\end{aligned}$$

The canopy tree is the limit $\boldsymbol{G}$ of the sequence $\boldsymbol{G_n}$. Its root will be random, depending on the probability of choosing a root on the $\boldsymbol{G_n}$. Generalizing the above observation, the probability of choosing a root at height k away from a leaf approaches $\frac{1}{2^{k+1}}$. Therefore the random rooted graph $\boldsymbol{G}$ which is a limit of $\boldsymbol{G_n}$ has the following probability distribution:

$$\Pr[\boldsymbol{G} = (T_C, x_n)] = \frac{1}{2^{n+1}}$$

□

It is perhaps surprisingly that the limit of truncated regular trees is not the infinite regular tree. This tree, as a matter of fact, is the limit of sequences of (some) regular graphs. We have seen this example in disguise: the cycles (2-regular graphs of increasing size) converge to the two-sided path (infinite 2-regular tree). The Benjamini-Schramm theory, once more, focuses on local properties of the graphs. Being a regular tree is a global property, and that was lost in the example of the canopy tree.

5 Conclusion: why the local convergence?

After seeing Benjamini-Schramm convergence in practice, we must return to the fundamental question: why is this graph limit theory valid?

First, several remarks throughout this exposition address the major feature of BS convergence: it is *local*. This happens because the limiting graph is determined by neighborhoods of the root. In turn, the theory works well for sparse graphs: the paths, the cycles, the trees all have Benjamini-Schramm limits. The sequence of complete graphs (which is dense) does not, for the degree of such sequence is not bounded.

Second, the Benjamini-Schramm space has nice topological properties. Because the space is separable by the set of finite graphs, we can use finite objects to efficiently approximate infinite networks for example. And, due to compactness, all sequences of bounded degree are guaranteed to have at least a convergent subsequence.

Third, this notion of convergence preserves several network statistics. Many of these statistics are "local" in the sense that they depend, mostly, on the size of neighborhoods. The clustering coefficient, a common measure of how tight communities are in social networks, is an example of network indicator which is estimable along the Benjamini-Schramm convergence under certain conditions. Similarly, so are measurements of degree assortativity. These and other results are found in section 2.5 of [6]. Matching numbers are another example of practically relevant indicator that is preserved under local convergence [12]. There are also many applications of this notion of convergence to Spectral Graph and Group Theory [13].

Finally, local convergence is being currently researched in the context of Machine Learning. Graph Neural Networks are learning algorithms often tasked with finding properties of nodes and edges, or predicting links. However, as posed in the introduction, many real-life networks are untractably large. Benjamini-Schramm convergence provides a framework for randomly sampling sections of the network in order to parallelize model training and achieve similar performance. Learning from local balls, as proven with the BS convergence, is as good as learning from the entire graph [14].

Therefore, for the particular class of sparse graphs, the local convergence has several advantages and applications in Mathematics and Computer Science.

References

[1] R. van der Hofstad, *Random Graphs and Complex Networks: Volume 1.* USA: Cambridge University Press, 1st ed., 2016.

[2] Y. Zhao, "Graph theory and additive combinatorics," 2019.

[3] R. Diestel, *Graph Theory.* Springer, August 2005.

[4] B. Bollobas, *Graph Theory: An Introductory Course.* New York: Springer, 1979.

[5] I. Benjamini and O. Schramm, "Recurrence of Distributional Limits of Finite Planar Graphs," *Electronic Journal of Probability*, vol. 6, no. none, pp. 1 – 13, 2001.

[6] R. van der Hofstad, "Random graphs and complex networks: Volume 2," 2021.

[7] C. Bordenav, "Lecture notes on random graphs and probabilistic combinatorial optimization," April 2016.

[8] N. Curien, "Random graphs: The local convergence point of view," September 2017.

[9] O. van Gaans, "Probability measures on metric spaces," 2003.

[10] S. Collective, “What is the limit of a sequence of graphs?? — benjamini-schramm convergence,” 2021.

[11] S. Collective, “The canopy tree: An example of a graph limit,” 2021.

[12] M. Abért, P. Csikvári, P. Frenkel, and G. Kun, “Matchings in benjamini-schramm convergent graph sequences,” 2021.

[13] S. Andraus, “Benjamini-schramm convergence and limiting eigenvalue density of random matrices,” 2018.

[14] T. Maehara and H. NT, “Learning on random balls is sufficient for estimating (some) graph parameters,” *CoRR*, vol. abs/2111.03317, 2021.