AN UPPER BOUND OF THE RECONSTRUCTION RADIUS OF RANDOM REGULAR GRAPHS

ALEXIS ZHOU AND CLAIRE ZHANG

Abstract. In the shotgun assembly problem, we are given rooted, local neighborhoods of a graph G, sampled from some distribution. We want to determine the minimum radius $R_\star$ of local neighborhoods allowing unique recovery of G, with high probability. In this paper, we study shotgun assembly of d-regular random graphs. We show that with high probability, $R_\star \leq \frac{\log n+\log\log n+\Delta}{2\log(d-1)}$ for an absolute constant Δ. Combined with [5], we establish an asymptotically tight bound: $R_\star \in \left[\frac{\log n+\log\log n-\Delta}{2\log(d-1)}, \frac{\log n+\log\log n+\Delta}{2\log(d-1)}\right]$ for an absolute constant Δ.

1. Introduction

We consider the shotgun assembly problem, proposed by Mossel and Ross [4]: there is an undirected graph $G = (V, E)$. For each $v \in V$, we are given the R-neighborhoods of v with only v labelled. Can G be assembled solely from these local neighborhoods?

This problem appears in various fields. In biology, DNA shotgun assembly is the reconstruction of a DNA sequence from smaller reads. This corresponds to seeing sections of a path graph, with vertices belonging to one of four colors. Researchers are interested in how large the reconstruction radius R must be, under various models of vertex labelling. Mossel and Ross [4] consider shotgun assembly of jigsaw puzzles given the neighborhoods of each piece. In machine learning, reconstructing a large neural networks from observed sub-networks is another instance of shotgun assembly.

Shotgun assembly is also a variant of the famous reconstruction conjecture in combinatorics, which states that a deterministic graph can be recovered uniquely from its list of vertex-deleted subgraphs [3, 2, 5]. In the present problem, the graph is random rather than worst-case, allowing us to take advantage of the structure of random regular graphs; however, compared to the setting of the reconstruction conjecture, we get access to smaller neighborhoods.

Bounds for this problem are known for various models of random graph. When G is an Erdős-Rényi random graph of constant average degree d, Mossel and Ross [4] showed that there are constants $0 < c_-(d) \leq c_+(d)$ such that, with high probability, assembly is possible if $R > c_+(d)\log n$, and not possible if $R < c_-(d)\log n$. When G is a random d-regular graph for $d \geq 3$, Bollobás [1] showed that assembly is possible for $R > (\frac{1}{2} + \epsilon)\frac{\log n}{\log d-1}$, for every $\epsilon > 0$. We also refer the reader to [4] for reconstruction bounds in other graph models, including the random jigsaw puzzle.

In the random regular setting, Mossel and Sun set out to establish a tighter bound:

Theorem 1.1. [5, Thm. 1] *Define*

$$R_- \equiv R_-(\Delta) \equiv \left\lfloor \frac{\log n + \log\log n - \Delta}{2\log(d-1)} \right\rfloor,$$
$$R_+ \equiv R_+(\Delta) \equiv \left\lceil \frac{\log n + \log\log n + \Delta}{2\log(d-1)} \right\rceil. \tag{1}$$

Let $G = (V, E)$ be a random d-regular graph on n vertices. Let $R_\star(G)$ be the minimal radius R required to assemble G from its list of rooted R-neighborhoods. Then there exists a positive absolute constant Δ such that

$$\lim_{n\to\infty} \mathbb{P}\left(R_-(\Delta) \leq R_\star(G) \leq R_+(\Delta)\right) = 1$$

for any fixed $d \geq 3$.

Date: July 31 2024.

2020 *Mathematics Subject Classification.* 05C80.

However, the proof of the upper bound contains errors. In this paper, we fix these errors, showing that Theorem 1.1 is indeed true.

Organization. The rest of this paper is structured as follows. In Section 2, we formally define the problem and establish key graph theoretic notions. In Section 3, we give a high-level overview of our proof of the upper bound on $R_\star$, outlining our solutions to two key issues in the proof of [5]. In Section 4 we present the lemma that fixes the first issue in [5], and introduce two w.h.p. properties that entail reconstructibility. The rest of the paper is devoted to proving the second property. In Section 5, we construct the coupling central to our proof. In Sections 6,7,8 we analyze the coupling, showing that w.h.p. we can mirror the growth of BFS's very successfully, while fixing [5]'s second issue. Section 9 concludes the coupling and proves the desired property.

2. Preliminaries

First, we formally define the problem.

Definition 2.1 (Shotgun assembly problem)**.** Let $G = (V, E)$. For a vertex $v \in V$, let $N_R(v)$ denote the subgraph of vertices in V that lie at graph distance $\leq R$ from v. Let $B_R(v)$ be $N_R(v)$ after removing edges $(u - w)$ such that $u, w \in N_R(v) \backslash N_{R-1}(v)$; in $B_R(v)$, vertices keep their original labels, and the position of v is specially marked. Let the R-neighborhood type of a vertex v be the isomorphism class $\mathcal{T}_R(v)$ of $B_R(v)$; in $\mathcal{T}_R(v)$, the vertices are no longer labelled, but the position of the root is still marked. We consider the question [4] of whether the graph G can be uniquely reconstructed, modulo global isomorphism, from its list $(\mathcal{T}_R(v))_{v\in V}$ of R-neighborhood types. The problem is to find $R_\star(G)$, the minimal radius R such that G can be uniquely reconstructed.

Per convention, we say an event E holds *with high probability (w.h.p.)* if $\mathbb{P}(E) = 1 - o_n(1)$. We will show that $R_\star(G)$ lies in a certain range w.h.p.. Towards this goal, we review several graph theoretic definitions originating from [5].

2.1. Graph theoretic definitions.

Definition 2.2 (Configuration model)**.** In the configuration model of n-vertex d-regular graphs, d *half-edges* are attached to n vertices. We let δv denote the set of half edges attached to a vertex v.

Let $\mathcal{G}_{n,d}$ be the set of n-vertex, d-regular graphs allowing self-loops and double edges. Let $\mathcal{G}'_{n,d}$ be the set of n-vertex, d-regular simple graphs (i.e. dis-allowing self-loops and double edges).

A perfect matching of the $\frac{nd}{2}$ half-edges creates a regular graph in $\mathcal{G}_{n,d}$. Accordingly, we can uniformly sample from $\mathcal{G}_{n,d}$ by choosing a perfect matching uniformly at random from $(nd-1)!!$ total options.

From this sampling procedure, the probability of sampling a simple graph is bounded away from 0 (depending on d) [6, Eq. 3]. If a property holds w.h.p. over the uniform distribution on $\mathcal{G}_{n,d}$, it also holds w.h.p. over the uniform distribution on $\mathcal{G}'_{n,d}$. Thus, we restrict our attention to graphs in $\mathcal{G}_{n,d}$, sampled with the configuration model.

For future use, we will additionally define N_R and B_R on half-edges: Let $G = (V, E)$. For a set of half-edges $\boldsymbol{s}$, let $\boldsymbol{t}$ be the set of incident vertices of half-edges in $\boldsymbol{s}$. Define $N_R(\boldsymbol{s})$ to be the subgraph of G induced on the set of vertices w, such that there exists a path of length $\leq R$ starting with some $g \in \boldsymbol{s}$ and ending at w that does not use any half-edge in $\bigcup_{v\in\boldsymbol{t}} \delta v \backslash \boldsymbol{s}$. Accordingly, define $B_R(\boldsymbol{s})$ by removing edges between distance-R vertices in $N_R(\boldsymbol{s})$, and marking the roots $\boldsymbol{t}$. From now on, we use this notation, and we refer to our previously defined $B_R(v)$ as $B_R(\delta v)$. We also define $\mathcal{T}_R(\boldsymbol{s})$, the rooted isomorphism type of $B_R(\boldsymbol{s})$.

Definition 2.3 (Breadth-first search exploration)**.** Given a graph $G = (V, E)$ and a list $\boldsymbol{s} = (g_1, \ldots, g_k)$ of distinct half-edges, the breadth-first search exploration (BFS) exploration of G started from $\boldsymbol{s}$ proceeds as follows. We will maintain a directed graph $\boldsymbol{H}_t$ of edges and vertices we have explored. We also maintain an ordered list F_t of frontier half-edges, which is the BFS queue. Initially, $\boldsymbol{H}_0$ is the graph with vertex set the union of incident vertices of g_i, and no edges; F_0 is $\boldsymbol{s}$. Let

$$\begin{aligned} \text{depth}(v) &\equiv 0 \quad \text{for all } v \in \boldsymbol{H}_0 \\ \text{par}(v) &\equiv v \quad \text{for all } v \in \boldsymbol{H}_0 \end{aligned}$$

At each time $t \geq 0$, as long as $F_t \neq \varnothing$, take the first half-edge g_t listed in F_t and reveal the half-edge h_t to which g_t is paired. Let u_t be the incident vertex of g_t and w_t be the incident vertex of h_t. Set $\boldsymbol{H}_{t+1}$ to be $\boldsymbol{H}_t$ together with a directed edge $u_t \to w_t$. If w_t was not already present in $\boldsymbol{H}_t$, then set

$$\operatorname{depth}(w_t) \equiv \operatorname{depth}(u_t) + 1$$
$$\operatorname{par}(w_t) \equiv u_t$$

and set F_{t+1} to be F_t with g_t removed and $\delta w_t \backslash \{h_t\}$ appended at the end:

$$F_{t+1} = (F_t \backslash \{g_t\}, \delta w_t \backslash \{h_t\})$$

(The half-edges incident to each vertex are ordered, so $\delta w_t \backslash \{h_t\}$ is an ordered list.) If w_t was already present in $\boldsymbol{H}_t$, then we term the creation of edge $(g_t - h_t)$ a *BFS collision*, and set F_{t+1} to be F_t with g_t, h_t removed. Accordingly, we define $\gamma(H)$ to be the set of collision events. After t steps, the number of unmatched half-edges remaining is $nd - 2t$. The process terminates upon reaching the first time t that $F_t = \varnothing$.

Remark 2.4. A BFS-exploration up to depth R terminates when we have finished exploring the edges from all depth-$(R-1)$ vertices, and found their depth R neighbors. Any edge between two depth-R vertices is not discovered in the exploration. Thus, $B_R(\boldsymbol{s})$ is exactly the graph explored by a depth-R BFS starting from $\boldsymbol{s}$.

Remark 2.5. Using the configuration model, we can simulate a BFS exploration starting from source half-edges $\boldsymbol{s} = s_1, \ldots, s_k$ without sampling the entire graph. We start with a set $[n]$ of isolated vertices, where each vertex v is equipped with an ordered list its half-edges. Let A_t denote the set of unmatched half-edges at time t, with $A_0 = [nd]$. Let $F_0 = (s_1, \ldots, s_k)$ list the initial half-edges in $\boldsymbol{s}$. At each time $t \geq 0$, as long as $F_t \neq \varnothing$, take the first half-edge g_t from F_t, and match it to a half-edge h_t which is sampled uniformly at random from $A_t \backslash \{g_t\}$. Set $A_{t+1} = A_t \backslash \{g_t, h_t\}$; and update F_t and $\boldsymbol{H}_t$ as in Definition 2.3. We refer to the sequence of random graphs $(\boldsymbol{H}_t)_{t \geq 0}$ as a simulated BFS. If nd is even, then it clearly has the same law as the BFS exploration in a graph G sampled from the configuration model $\mathbb{P}_{n,d}$.

We use the notion of cycle structure to encode neighborhood structures in a more compact way.

Definition 2.6 (cycle structure)**.** Let the *cycle structure* of a directed graph H rooted at r be an undirected graph $\mathcal{C}(H)$, consisting of

- all cycles in the undirected version of H
- all (directed) edges $u \to v$ in H such that $v \in \mathcal{C}(H)$.

The full neighborhood of a vertex can be inferred from its cycle structure because we know the degree of each vertex. Dangling trees do not contribute to cycle structure.

Lemma 2.7. *When $H = H_t$ is the result of a t-step BFS-exploration from one source r, the evolution of $\mathcal{C}_t \equiv \mathcal{C}(H_t)$ can be described as follows. At the beginning of the exploration, $\mathcal{C}_0 = \varnothing$. Let*

$$\operatorname{par}^\ell(x) = \overbrace{\operatorname{par}(\cdots \operatorname{par}}^{\ell}(x) \cdots)$$

be the ℓ-th ancestor of vertex x. For time $t \geq 0$, if $g_t \to h_t$ is a collision, set

$$\mathcal{C}_t = \mathcal{C}_{t-1} \cup (u_t \to w_t) \cup \bigcup_{\ell \geq 0} \left(\operatorname{par}^{\ell+1}(u_t) \to \operatorname{par}^\ell(u_t)\right) \cup \bigcup_{\ell \geq 0} \left(\operatorname{par}^{\ell+1}(w_t) \to \operatorname{par}^\ell(w_t)\right).$$

That is, append the newly formed edge and the ancestor paths of u_t and w_t (see Figure 1). Otherwise, set $\mathcal{C}_t = \mathcal{C}_{t-1}$.

Proof. We show this process generates the cycle structure. First we show that $H_t \backslash \mathcal{C}_t$ is a tree in the quotient graph $H_t \backslash \mathcal{C}_t$, rooted at the super-node $\mathcal{C}_t$. In a tree, the addition of a tree edge doesn't change its cycle structure; the addition of a non-tree edge introduces a single cycle plus an upward path to the cycle structure.

We proceed with induction: $\mathcal{C}_0 = \mathcal{C}(H_0)$. $G' \subset G$ implies $\mathcal{C}(G') \subset \mathcal{C}(G)$, so $\mathcal{C}_{t-1} \subset \mathcal{C}_t$. The addition of $(g_t \to h_t)$ to a tree (in the quotient graph) creates at most one cycle, and the only paths to u_t or w_t are tree edges. Thus, $\mathcal{C}_t \backslash \mathcal{C}_{t-1}$ is precisely $(g_t \to h_t)$ plus the ancestor path from u_t to $\mathcal{C}_{t-1}$ plus the ancestor path from w_t to $\mathcal{C}_{t-1}$. □

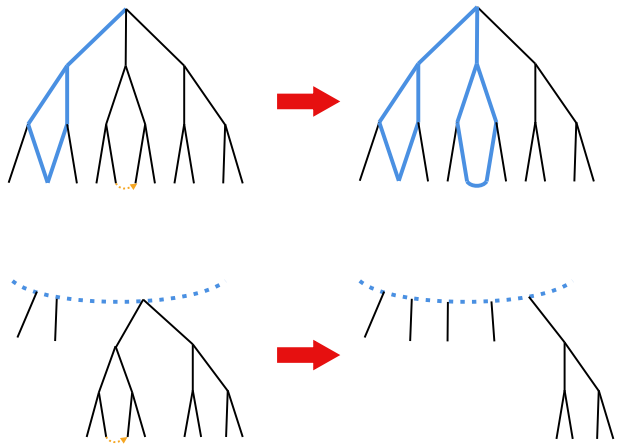

FIGURE 1. Two equivalent views of a collision (yellow) affecting cycle structure (blue).

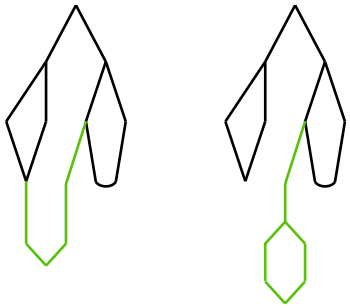

FIGURE 2. *add* (left: case 1, right: case 2).

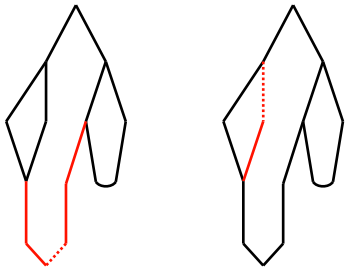

FIGURE 3. *delete* the dashed edge, removing all red edges.

Definition 2.8. Let H be a directed graph and root set $\boldsymbol{t}$ be a set of vertices. Define a $\boldsymbol{t}$-cycle to be the preimage (in H) of a cycle in the undirected version of $H \backslash \boldsymbol{t}$. The cycle structure of H given $\boldsymbol{t}$ is an undirected graph $\mathcal{C}(H)$, consisting of

- all $\boldsymbol{t}$-cycles
- all (directed) edges $u \to v$ in H such that $v \in \mathcal{C}(H)$.

Remark 2.9. When $|\boldsymbol{t}| > 1$, a cycle in H must be an $\boldsymbol{t}$-cycle, but not always vice versa. Our previous definition of cycle structure agrees with the $\boldsymbol{t} = \{r\}$ case. $\mathcal{C}(H)$ with root set $\boldsymbol{t}$ is a union of the cycle structures of each root in $\boldsymbol{t}$, plus paths that connect different root vertices and their cycle structures.

Let $\boldsymbol{s}$ be the set of source half-edges, $\mathcal{C}(B_R(\boldsymbol{s}))$ denotes the cycle structure of $B_R(\boldsymbol{s})$, with root set being the set of incident vertices of half-edges in $\boldsymbol{s}$.

We quantify how far two cycle structures are from each other with a custom distance metric, defined over cycle structures with one root vertex.

Definition 2.10 (distance between cycle structures)**.** Given a rooted cycle structure $\mathcal{C}$, we define the following operations:

A *delete* operation cuts an edge e in $\mathcal{C}$ that preserves the cycle structure's connectivity, then recursively removes the leaves (see Figure 3).

An *add* operation is the reverse of a delete operation. There are two cases (see Figure 2):

(1) Choose distinct vertices w, w' in $\mathcal{C}$ and an integer $l > 0$. Attach a path of length l between w and w' that doesn't intersect with any other vertices in $\mathcal{C}$.
(2) Choose a vertex w in $\mathcal{C}$ and two integers $l_1 \geq 0, l_2 > 0$. Attach a path of length l_1 to w, with a cycle of length l_2 at the end of the path, such neither the path nor cycle intersects with any other vertex in $\mathcal{C}$.

Given two rooted cycle structures $\mathcal{C}, \mathcal{C}'$, we define $dist\,(\mathcal{C}, \mathcal{C}')$ to be the minimum number of *add* and *delete* operations required to transform $\mathcal{C}$ into $\mathcal{C}'$.[1]

By definition $dist(\mathcal{C}, \mathcal{C}) = 0$ and the triangle inequality holds. If $dist\,(\mathcal{C}, \mathcal{C}') = m$, then we can trace a sequence of m operations to go from $\mathcal{C}$ to $\mathcal{C}'$, where each operation is reversible by construction. Hence $dist\,(\mathcal{C}', \mathcal{C}) = m$ and $dist$ is indeed a metric.

2.2. **A priori bounds on random regular graphs.** We will use the following well-known form of the Chernoff bound whenever we argue by stochastic dominance: if X is a binomial random variable with mean μ, then for all $t \geq 1$,

$$\mathbb{P}\,(X \geq t\mu) \leq \exp\{-t\mu \log\,(t/e)\}.$$

In addition, the following bounds relating to the cycle structure of random regular graphs will be useful:

Lemma 2.11 (cycle packing bound, [5, Lem. 3.8])**.** *Suppose for $R \geq 0$ that $\mathcal{C}$ is the depth-R cycle structure of a vertex v in a d-regular graph. Then, as long as $\mathcal{C}$ is non-empty,*

$$|E(\mathcal{C})| \leq 2\gamma(\mathcal{C})\left(R - \log_{d-1}\gamma(\mathcal{C}) + 1 + o_d(1)\right).$$

Lemma 2.11 applies to all R, while the proceeding lemmas apply to any R at most

$$R_{\max} \equiv \frac{\log n + 2\log\log n}{2\log(d-1)}.$$

Note that $R_{\max} \geq R_+(\Delta)$ for sufficiently large n.

Lemma 2.12 (total number of collisions, [5, Lem. 3.9])**.** *Let $G = (V, E)$ be a random d-regular graph on n vertices. Let $\mathcal{C} = \mathcal{C}(B_R(\boldsymbol{s}))$ where $\boldsymbol{s}$ is the list of source half-edges.*

Let k be the number of distinct incident vertices of the half-edges in $\boldsymbol{s}$. If k is constant (not depending on n) and $R \leq R_{\max}$, then

$$\mathbb{P}\left(\gamma(\mathcal{C}) \geq (2ek\log n)^2\right) \leq \exp\left\{-(\log n)^2\right\}$$

for sufficiently large n.

Lemma 2.13 (few shallow collisions, [5, Lem. 3.10])**.** *Let $G = (V, E)$ be a random d-regular graph on n vertices. Let $\mathcal{C} = \mathcal{C}(B_R(\boldsymbol{s}))$. Let k be the number of distinct incident vertices among $\boldsymbol{s}$.*

If k is constant and $R \leq \frac{1-\epsilon}{2}\log_{d-1} n$ with $\frac{1}{\log n} \ll \epsilon < 1$, then

$$\mathbb{P}(\gamma(C) \geq \rho/\epsilon) \leq n^{-\rho + o_n(1)}$$

for any positive constant ρ.

Lemma 2.14 (few short cycles, [5, Lem. 3.11])**.** *Let $G = (V, E)$ be a random d-regular graph on n vertices. Let $\mathcal{C} = \mathcal{C}(B_R(\boldsymbol{s}))$ where $|\boldsymbol{s}|$ is a constant and $R \leq R_{\max}$; and let $\mathcal{C}_\circ \subset \mathcal{C}$ denote the cycle structure induced by cycles in $B_R(\boldsymbol{s})$ with length at most $(1-\epsilon)\log_{d-1} n$. If $\frac{\log\log n}{\log n} \ll \epsilon < 1$, then*

$$\mathbb{P}(\gamma(\mathcal{C}_\circ) \geq 2\rho/\epsilon) \leq n^{-\rho + o_n(1)}$$

for any positive constant ρ.

3. Technical overview

The proof of the upper bound uses an insight from [4], that if a graph's R-neighborhoods are non-isomorphic, then the graph can be reconstructed from the $(R+1)$-neighborhoods. This is because, given the $(R+1)$-neighborhood of any vertex v, we can identify the neighbors of v through their R neighborhoods. So, we may reconstruct the graph by starting at an arbitrary vertex v, identifying its neighbors, and then iteratively the neighbors of these vertices, until we propagate through the whole graph. Thus, it suffices to prove that all $R = (R_+(\Delta) - 1)$-neighborhoods are distinct. Proving this claim is also the strategy taken in [5].

If we imagine for a moment that all R-neighborhoods of G are independent random variables, then this distinctness follows easily. Indeed, one can calculate that the probability of realizing any R-neighborhood

[1] Note that this is not the edit distance, but one that accounts for the number of cycles. It is not a homotopy, as lengths of paths and cycles matter.

isomorphism class (identified by its cycle structure) is $o(n^{-2})$ (see [5, Proposition 4.1]). Thus for any u, v the probability $\mathbb{P}(B_R(u) \cong B_R(v)) = o(n^{-2})$, and taking a union bound over pairs of vertices u, v implies all neighborhoods are pairwise distinct with high probability.

However, in reality the R-neighborhoods of G are not independent. First, if $u, v \in G$ are close to each other, their R-neighborhoods have large overlap and are clearly not independent. Second, we will see that for our choice of R, the R-neighborhoods of even typical pairs of vertices u, v will intersect, and therefore not be independent.

Mossel and Sun [5] introduce two main techniques to address the non-independence of neighborhoods. To address the non-independence of $B_R(u)$ and $B_R(v)$ for nearby u, v, they argue that the BFS explorations around u, v differ "in some direction." To address that the BFS explorations around even far away vertices u, v collide with high probability, they introduce a coupling argument comparing $B_R(u)$ and $B_R(v)$ to a pair of independent BFS explorations $B_R(x)$ and $B_R(y)$.

However, the implementations of these two ideas contain two main issues, which are fixed in the present paper and summarized below.

3.1. **Issue 1: deducing non-isomorphism of neighborhoods from non-isomorphism of directed neighborhoods.** To address that the neighborhoods $B_R(u)$ and $B_R(v)$ can have large overlap if u, v are nearby in G, [5] aims to argue that $B_R(u)$ and $B_R(v)$ differ "in some direction." Formally, they argue via the following two-step strategy:

- First, they show that for each $u, v \in G$, there are sets of "good directions" $\boldsymbol{u} \subseteq \delta u, \boldsymbol{v} \subseteq \delta v$ of size $|\boldsymbol{u}| = |\boldsymbol{v}| = d - 2$, such that (among other conditions) the directed BFS's of $B_R(\boldsymbol{u}), B_R(\boldsymbol{v})$ do not intersect up to depth $L_\circ = \frac{1}{16} \log_{d-1} n \approx R/8$.
- Then, they show that *conditional on these properties*, w.h.p. $B_R(\boldsymbol{u}) \ncong B_R(\boldsymbol{v})$. This is achieved by the coupling argument described in the next subsection.

The proof of the second step requires that $B_R(\boldsymbol{u}), B_R(\boldsymbol{v})$ do not intersect up to depth $L_\circ$. This is because otherwise, $B_R(\boldsymbol{u})$ and $B_R(\boldsymbol{v})$ will share too much structure, and the part where they grow more or less "independently" is too small for the coupling method to work. This limitation applies to our proofs as well and guides how we design the arguments below.

Unfortunately, the result of the above argument, that $B_R(\boldsymbol{u}) \ncong B_R(\boldsymbol{v})$, does not imply that $B_R(\delta u) \ncong B_R(\delta v)$, as a putative isomorphism $\varphi : B_R(\delta u) \to B_R(\delta v)$ does not necessarily map $\boldsymbol{u}$ to $\boldsymbol{v}$.

Our proof finds a way to link non-isomorphism of directed neighborhoods to non-isomorphism of neighborhoods. Here we present a simplified version of this argument, which holds for $d \geq 5$. The general $d \geq 3$ case follows from a more involved combinatorial argument, which is carried out in subsection 9.1.

We will show that if a graph G satisfies the following two properties, the $(R+1)$-neighborhoods of G are pairwise non-isomorphic.

(1) For all $u, v \in G$, there exist directions $\boldsymbol{u} \subseteq \delta u$, $\boldsymbol{v} \subseteq \delta v$ of size $|\boldsymbol{u}| = |\boldsymbol{v}| = d - 2$, such that $B_{L_\circ+1}(\boldsymbol{u})$ and $B_{L_\circ+1}(\boldsymbol{v})$ are non-intersecting complete trees.
(2) For all $u, v \in G$, and *all* $\boldsymbol{u} \subseteq \delta u, \boldsymbol{v} \subseteq \delta v$ of size $|\boldsymbol{u}| = |\boldsymbol{v}| = d - 2$, either (a) or (b) holds, where
   (a) the neighborhoods $B_{L_\circ}(\boldsymbol{u})$ and $B_{L_\circ}(\boldsymbol{v})$ are *not* disjoint trees
   (b) the neighborhoods $B_R(\boldsymbol{u})$ and $B_R(\boldsymbol{v})$ are not isomorphic.

Property (1) holds with high probability by a union bound applied to Lemma 2.13. We will outline the proof of property (2) in the next subsection.

We next explain why these two properties imply that $B_R(u) \ncong B_R(v)$ for all $u, v \in G$. Suppose for contradiction that there is an isomorphism $\varphi : B_{R+1}(u) \mapsto B_{R+1}(v)$, and let $\boldsymbol{u} \subseteq \delta u$, $\boldsymbol{v} \subseteq \delta v$ be the directions given by property (1), so that $B_{L_\circ+1}(\boldsymbol{u})$ and $B_{L_\circ+1}(\boldsymbol{v})$ are non-intersecting trees; see Figure 4.

Since $\boldsymbol{u}$ and $\boldsymbol{v}$ each exclude only 2 half-edges of δu and δv, for $d \geq 5 > 2 + 2$ there exists $e \in \boldsymbol{u}$ such that $B_{L_\circ+1}(e)$ and $B_{L_\circ+1}(\varphi(e))$ are non-intersecting trees. By property (2) (applied to the other endpoints of e, $\varphi(e)$ and any $d-2$ of the $d-1$ descendant directions), $B_{R+1}(e)$ and $B_{R+1}(\varphi(e))$ are not isomorphic, which contradicts $\varphi(B_{R+1}(u)) = B_{R+1}(v)$.

3.2. **Issue 2: proving (2) through a coupling.** It remains to prove (2). If $B_R(\boldsymbol{u}), B_R(\boldsymbol{v})$ are not disjoint trees in the upper $L_\circ$-levels, the (a) is automatically satisfied. So we only look at instances in the joint

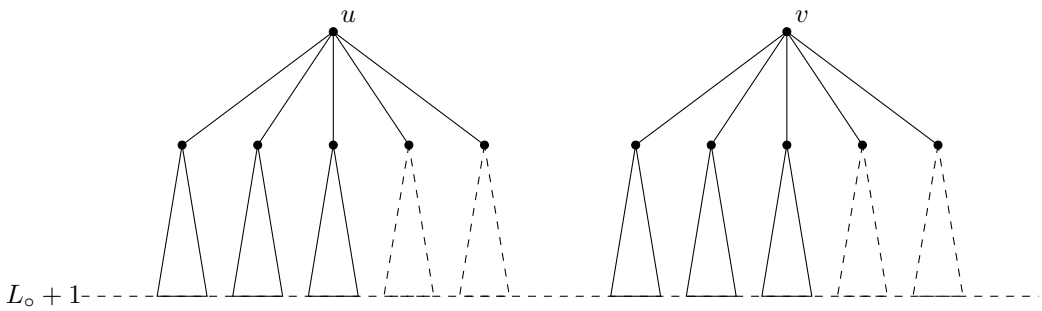

FIGURE 4. Solid triangles in represent explorations that are trees and do not intersect the rest of the graph, whereas dotted triangles may intersect $B_R(v)$ or not be trees. If $d \geq 5$, any putative isomorphism $\varphi : B_R(u) \to B_R(v)$ must map at least one solid tree in $B_R(u)$ to one solid tree in $B_R(v)$. By property (2) this is not possible, and therefore $B_R(u) \ncong B_R(v)$.

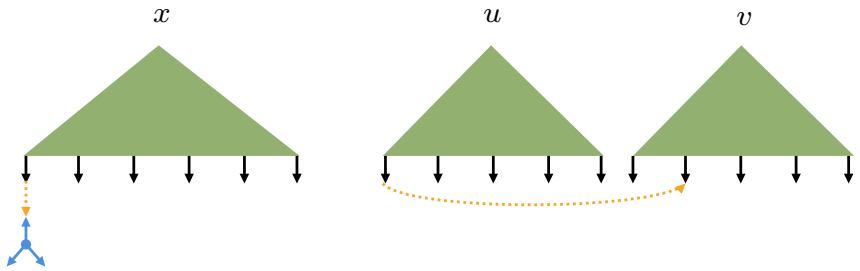

FIGURE 5. When a "crossing" (yellow dashed edge) occurs in the $\boldsymbol{u}$-exploration, we couple it with the $\boldsymbol{x}$-exploration finding a fresh vertex.

distribution where we have disjoint trees. In these scenarios, large portions of the two graphs are still "independent", and therefore unlikely to coincide.

Following [5], we construct a coupling of the u, v BFS co-evolution with two *independent* BFS's: $\boldsymbol{x}$ and $\boldsymbol{y}$ which live in two different graphs. The growth of the $\boldsymbol{u}$-BFS will imitate $\boldsymbol{x}$, and growth of the $\boldsymbol{v}$-BFS will imitate $\boldsymbol{y}$. So in each instance of the joint $(\boldsymbol{x}, \boldsymbol{uv}, \boldsymbol{y})$ distribution, $B_R(\boldsymbol{u})$ and $B_R(\boldsymbol{x})$ (as well as $B_R(\boldsymbol{v})$ and $B_R(\boldsymbol{y})$) will agree as much as possible.

It suffices to prove that w.h.p. $B_R(\boldsymbol{u})$ and $B_R(\boldsymbol{x})$ are $O(1)$ in cycle structure distance, and that the same holds for $B_R(\boldsymbol{y})$ and $B_R(\boldsymbol{v})$. [5] shows that $dist(B_R(\boldsymbol{x}), B_R(\boldsymbol{y})) > \frac{\log n}{\log \log n}$ w.h.p., due to their independence. $O(1)$ distance in the coupling implies that $dist(B_R(\boldsymbol{u}), B_R(\boldsymbol{v})) > \frac{\log n}{\log \log n} - O(1)$ w.h.p.. Therefore, the probability that $B_R(\boldsymbol{u}), B_R(\boldsymbol{v})$ are isomorphic with zero distance is very small ($o(n^{-2})$). And a union bound over pairs u, v gives w.h.p. reconstructibility.

Now, it remains to show $dist(B_R(\boldsymbol{x}), B_R(\boldsymbol{u})) = O(1)$ whp. We bound the distance by deleting cycles in $B_R(\boldsymbol{u})$ and $B_R(\boldsymbol{x})$ where they do not agree, until we are left with the part shared between them. The distance is upper bounded by the number of deletions we make in both graphs combined. The proof of this bound in [5] had an error, and fixing this error is the main contribution of the present paper.

The presence of $\boldsymbol{v}$ in the same graph affects how the $\boldsymbol{u}$-exploration evolves. So the laws of $B_R(\boldsymbol{x})$ and $B_R(\boldsymbol{u})$ are not the same, and we cannot couple $\boldsymbol{x}$ and $\boldsymbol{x}$ perfectly. A coupling error is a BFS step where the co-evolved $\boldsymbol{x}$- and $\boldsymbol{u}$-explorations fail to agree (ref subsection 5.3). Errors come in many kinds, but the most prevalent by far is where an edge traverses the part explored by $\boldsymbol{u}$ and the part explored by $\boldsymbol{v}$. We call such edge a "crossing". Crossings occur $\log n$ many times, until we end the BFS at level R. Meanwhile, all remaining coupling errors occur $O(1)$ many times, and can be undone with $O(1)$ operations.

A crossing edge, at depth l in the $\boldsymbol{u}$-neighborhood, absorbs into $B_R(\boldsymbol{u})$ a neighborhood of radius $R-l-1$ in the $\boldsymbol{v}$-graph. This neighborhood is determined by $\boldsymbol{v}$, and hence controlled by the $\boldsymbol{y}$ exploration, which the $\boldsymbol{x}$-exploration has no hope to resemble. Thus, when a crossing occurs, we let the $\boldsymbol{x}$- (or $\boldsymbol{y}$-, if the edge goes from the $\boldsymbol{v}$-graph to the $\boldsymbol{u}$-graph) exploration find a fresh vertex instead (see Figure 5). In order to remove

the disagreeing parts of $B_R(\boldsymbol{x})$ and $B_R(\boldsymbol{u})$, we need to remove all cycles involving the crossings or crossing neighborhoods.

The crux of our proof is that despite there being $\log n$ crossings, their effect on the cycle structure of $B_R(\boldsymbol{u})$ can be undone by cutting out $O(1)$ cycles.

We will first argue that O(1) operations suffice to delete all within-neighborhood cycles, thanks to our $L_\circ$ non-intersection condition. The depth of a crossing $l \geq L_\circ$. So the neighborhood radius $R - l - 1 \leq \frac{7.1}{8} R$ is relatively small, so all within-neighborhood cycles are short. And there are not many short cycles (ref subsesction 7.2). Then, we argue that any pair of intersecting neighborhoods can also be disconnected with $O(1)$ operations (ref subsection 7.3), again because crossing neighborhoods are shallow in radius. This alone does not suffice because there are $\log_{d-1} n$ neighborhoods. However, most of the neighborhoods, *especially those coming from crossings from greater depths*, do not intersect with any other neighborhood. We show this by defining *upper* and *lower* crossings. The upper crossings create *large and probably intersecting neighborhoods*, but there are *only* $O(1)$ *of them.* All lower crossings create *very small neighborhoods that appear in random positions in the v-graph, and hence they rarely intersect.* Still, there are $\log n$ *of them.* We can w.h.p. upper bound the number of intersecting neighborhood pairs by $O(1)$. Hence, $O(1)$ deletions suffice to erase the $\boldsymbol{v}$-graph from the cycle structure of $B_R(\boldsymbol{u})$.

4. Distinguishability of two neighborhoods in the same graph: overview

Throughout the rest of the paper, we require $R \leq R_{\max}$. Our goal is to prove the upper bound of the reconstruction radius, namely: given the $R_+(\Delta)$-depth neighborhoods of every vertex, we can reconstruct the random d-regular neighborhood w.h.p. as $n \to \infty$.

Definition 4.1. Throughout what follows we denote $L_\circ \equiv \frac{1}{16} \log_{d-1} n$.

Let $\mathcal{T}_{dir,L_\circ}$ denote the rooted, depth-$L_\circ$, directed d-regular tree, where the root has degree $d-2$, every other vertex above depth $L_\circ$ has degree d, and the depth-$L_\circ$ vertices have degree 1.

4.1. Proof overview.

We prove probability $1 - o_n(1)$ reconstructibility by the following lemma:

Lemma 4.2. *For any fixed d-regular $G = (V, E)$ that satisfies the following:*

1. $|\gamma(\mathcal{C}(B_{L_\circ+1}(\delta u \cup \delta v)))| \leq 2$ *for any $u \neq v \in V$.*
2. *For any $u \neq v \in V$ and any directions $\boldsymbol{u} \subset \delta u, \boldsymbol{v} \subset \delta v$, $|\boldsymbol{u}| = |\boldsymbol{v}| = d-2$, either (a) does not hold, or (b) holds. Here,*
   (a) $B_{L_\circ}(\boldsymbol{u}) \cap B_{L_\circ}(\boldsymbol{v}) = \varnothing$ *and* $B_{L_\circ}(\boldsymbol{u}) \cong B_{L_\circ}(\boldsymbol{v}) \cong \mathcal{T}_{dir,L_\circ}$
   (b) $B_R(\boldsymbol{u}) \ncong B_R(\boldsymbol{v})$

We have $B_{R+1}(u) \ncong B_{R+1}(v)$ for any $u \neq v \in G$. Hence, once we are given all the depth-$R+2$ neighborhoods of G, we can reconstruct G.

We defer the proof of Lemma 4.2 to subsection 9.1. If (1) holds w.h.p. over all random $d-$regular graphs, and (2) holds w.h.p. with $R = R_+(\Delta) - 2$, then by union bound over (1) and (2), a d-random regular graph is reconstructible given all neighborhoods of depth $R_+(\Delta)$.

Hence, to establish the upper bound of the reconstruction radius, it suffices to prove the following two lemmas:

Lemma 4.3. *W.h.p. over all random d-regular graphs, item (1) in Lemma 4.2 holds.*

Proof. In Lemma 2.13, let $\boldsymbol{s} = \{\delta u, \delta v\}$, $k = 2$, $\rho = 2.1$, $\epsilon = 1 - \frac{2(L_\circ+1)}{\log_{d-1} n} \geq \frac{7}{9}$, then $\mathbb{P}(\gamma(\mathcal{C}(B_{L_\circ+1}(\delta u \cup \delta v))) \geq 2.7) \leq \mathbb{P}(\gamma(\mathcal{C}(B_{L_\circ+1}(\delta u \cup \delta v))) \geq \rho/\epsilon) \leq n^{-2.1+o_n(1)}$. By union bound over all $u \neq v$, $\gamma(\mathcal{C}(B_{L_\circ+1}(\delta u \cup \delta v))) \leq 2$ holds for all $u \neq v$ with probability $1 - n^{-0.1+o_n(1)} = 1 - o_n(1)$. □

Lemma 4.4. *Given vertices $u \neq v \in V$, and directions $\boldsymbol{u} \subset \delta u, \boldsymbol{v} \subset \delta v$, $|\boldsymbol{u}| = |\boldsymbol{v}| = d-2$, let event $\mathcal{E}_0(\boldsymbol{u}, \boldsymbol{v})$ be the event defined by (2) in Lemma 4.2, event $\mathcal{E}$ be the event defined in Lemma 4.5 below, then*

$$\mathbb{P}(\mathcal{E} \text{ and not } \mathcal{E}_0(\boldsymbol{u}, \boldsymbol{v})) \leq o(n^{-2})$$

where the probability is over all random d-regular graphs on n vertices.

If Lemma 4.4 is true, then $\mathbb{P}(\text{not } \mathcal{E}_0(\boldsymbol{u}, \boldsymbol{v}) \text{ for any } \boldsymbol{u}, \boldsymbol{v}) \leq \mathbb{P}(\text{not } \mathcal{E}) + \mathbb{P}(\mathcal{E} \text{ and not } \mathcal{E}_0(\boldsymbol{u}, \boldsymbol{v}) \text{ for any } \boldsymbol{u}, \boldsymbol{v}) \leq o_n(1) + \Theta(n^2)\binom{d}{d-2}^2 \cdot o(n^{-2}) = o_n(1)$, where the second inequality follows from Lemma 4.4 and a union bound

over all $u \neq v, \boldsymbol{u}, \boldsymbol{v}$. So we know item (2) of Lemma 4.4 holds w.h.p.. Combining with Lemma 4.3, we proved that reconstruction is possible w.h.p..

The following sections are all devoted to proving Lemma 4.4, through a coupling on the directed neighborhoods.

4.2. **Probability of realizing a particular directed neighborhood structure.** Define $\mathbf{\Omega}_{dir,R}$ to be the set of all neighborhood isomorphism types $\mathcal{T}_R = \mathcal{T}_R(\boldsymbol{v})$ which can arise from a d-regular graph with $|\boldsymbol{v}| = d-2$, such that

$$|E(\mathcal{T}_R)| \geq (d-1)^R \left(\frac{d-2}{d-1} - \frac{11}{\log n}\right).$$

$\mathbf{\Omega}_{dir,R}$ represents a set of "reasonable" neighborhood structures, which do not have too many collisions. In order for $\mathcal{E}_0(\boldsymbol{u}, \boldsymbol{v})$ to hold with $1 - o(n^{-2})$ probability, we need neighborhood structures similar to $B_R(\boldsymbol{u})$ to be "reasonable", hence the $\mathcal{E}$ in Lemma 4.4.

Lemma 4.5. *W.h.p. over all random d-regular graphs $G = (V, E)$ on n vertices, the following holds: for any $u \in V$, $\boldsymbol{u}$ with $|\boldsymbol{u}| = d-2$, and directed type $\mathcal{Q}_R$ ($\mathcal{Q}$ is not part of G), if $\mathcal{Q}_{L_\circ} \cong B_{L_\circ}(\boldsymbol{u}) \cong \mathcal{T}_{dir,L_\circ}$ and $dist(\mathcal{C}(\mathcal{Q}_R), \mathcal{C}(B_R(\boldsymbol{u}))) \leq (\log n)^2$, then $\mathcal{Q}_R \in \mathbf{\Omega}_{dir,R}$.*

Proof. This follows from combining Lemma 4.5 (after taking union bound on all vertices), Lemma 5.6, Corollary 5.7 in [5]. □

Lemma 4.6 (Corollary 5.7 of [5])**.** *For any positive constant ρ, there exists $\Delta \equiv \Delta(\rho)$ sufficiently large so that for $R \geq R_+(\Delta)$, and any fixed $\boldsymbol{v} \subset \delta v$ with $|\boldsymbol{v}| = d-2$,*

$$\mathbb{P}(\mathcal{T}_R(\boldsymbol{v}) = \mathcal{T}_R) \ll n^{-\rho} \quad \text{for all } \mathcal{T}_R \in \mathbf{\Omega}_{\mathrm{dir},R}.$$

Remark 4.7. Lemma 4.5 says when a neighborhood structure will be "reasonable". Lemma 4.6 is the analogue of Proposition 4.1 in [5]: the chance of realizing *any* reasonable directed neighborhood is small. Also, Lemma 4.6 is the only place in Section 4 where we need R to be large enough (i.e. $R \geq R_+(\Delta)$).

5. Distinguishability: a coupling to approximate independence

If either $B_R(\boldsymbol{u})$ and $B_R(\boldsymbol{v})$ are not are not disjoint trees up to depth L_0, $\mathcal{E}_0(\boldsymbol{u}, \boldsymbol{v})$ is automatically satisfied. When we have disjoint trees, we will BFS simulate the (much more complicated) depth-R neighborhoods.

Given Lemma 4.6, if $B_R(\boldsymbol{u}), B_R(\boldsymbol{v})$ were independent, their probability of being isomorphic is small. When $B_R(\boldsymbol{u}), B_R(\boldsymbol{v})$ reside in the same graph, they affect each other: probability of each step being a collision changes, and the two neighborhoods could intersect.

5.1. **An overview of the coupling and its analysis.** To capture the observation that the majority of probabilities and edges of $B_R(\boldsymbol{u})$, $B_R(\boldsymbol{v})$ are still "free", we couple the step-by-step growth of $B_R(\boldsymbol{u})$ with the growth of neighborhood $B_R(\boldsymbol{x})$, and $B_R(\boldsymbol{v})$ with $B_R(\boldsymbol{y})$, where $\boldsymbol{x}, (\boldsymbol{u}, \boldsymbol{v}), \boldsymbol{y}$ exist in three separate graphs. The neighborhoods $B_R(\boldsymbol{u})$ and $B_R(\boldsymbol{v})$ are correlated because u, v belong to the same graph. We will require the neighborhoods $B_R(\boldsymbol{x})$ and $B_R(\boldsymbol{y})$ to be independent. The coupling will attempt to make the growth of $B_R(\boldsymbol{u})$ imitate that of $B_R(\boldsymbol{x})$ as much as possible, and similarly for $B_R(\boldsymbol{v})$ and $B_R(\boldsymbol{y})$. $K_t(\boldsymbol{x}) \cong K_t(\boldsymbol{u})$ is the "good" part of the $B_R(\boldsymbol{x})$ and $B_R(\boldsymbol{u})$ that can be matched, due to the imitation, and w.h.p. a significant part.

The end result would be a joint distribution of $B_R(\boldsymbol{x}), (B_R(\boldsymbol{u}), B_R(\boldsymbol{v})), B_R(\boldsymbol{y})$. The marginal distributions of $B_R(\boldsymbol{x})$ and $B_R(\boldsymbol{y})$ are both the distribution of a depth-R $(d-2)$-directed neighborhood in a random d-regular graph on n vertices. The marginal distribution of $(B_R(\boldsymbol{u}), B_R(\boldsymbol{v}))$ is that of two neighborhoods in the same random d-regular graph on n vertices. The marginal of $B_R(\boldsymbol{x}), B_R(\boldsymbol{y})$ is independent.

For each instance of $B_R(\boldsymbol{x}), (B_R(\boldsymbol{u}), B_R(\boldsymbol{v})), B_R(\boldsymbol{y})$ in the joint distribution, we first check if (a) is satisfied. If not, $\mathcal{E}_0(\boldsymbol{u}, \boldsymbol{v})$ automatically holds.

If so, we will bound the distance between $\mathcal{C}(B_R(\boldsymbol{x}))$ and $\mathcal{C}(B_R(\boldsymbol{u}))$ by cutting edges to restore them both to the shared "good" part. This distance is small w.h.p., so is $dist(\mathcal{C}(B_R(\boldsymbol{v})), \mathcal{C}(B_R(\boldsymbol{y})))$. Not so many neighborhoods $B_R(\boldsymbol{x})$ are within a constant distance of $B_R(\boldsymbol{y})$, so by the previous section, $dist(\mathcal{C}(B_R(\boldsymbol{x})), \mathcal{C}(B_R(\boldsymbol{y})))$ is large w.h.p.. As a result, $dist(\mathcal{C}(B_R(\boldsymbol{u})), \mathcal{C}(B_R(\boldsymbol{v}))) > 0$ w.h.p., and $\mathcal{E}_0(\boldsymbol{u}, \boldsymbol{v})$ holds.

5.2. **Details of the coupling: definitions.** To achieve the aforementioned joint distribution, we define a coupling on three BFS explorations.

x is a vertex in G_x. $u \neq v$ are vertices in G_{uv}. y is a vertex in G_y. At time 0 these graphs are empty: each vertex has d-half-edges attached, but there is no edge, as we will match half-edges to form edges step-by-step. These are three completely separate n-vertex d-regular graphs (no two graphs share any vertex). The directions satisfy $|\boldsymbol{x}| = |\boldsymbol{u}| = |\boldsymbol{v}| = |\boldsymbol{y}| = d - 2$.

In step t, one or two edge(s) form by matching half-edges. The edge(s) formed can live in any of the three graphs. The subscript $_t$ denotes the time before the step-t edge(s) form, but after the step-$(t-1)$ edge(s) form. The exploration terminates when we have found $B_R(\boldsymbol{x}), B_R(\boldsymbol{u}), B_R(\boldsymbol{v}), B_R(\boldsymbol{y})$. We let the last edge to be formed in the $t(R)$-th step.

Throughout time ($0 \leq t \leq t(R)$), we will keep track of the following entities:

- F_t is the *frontier*, an ordered sequence of certain unmatched half-edges possibly from all three graphs, waiting to be matched in order. We will add and remove half-edges from F_t as we would do in a BFS (ref Definition 2.3).
- g_t, the half-edge matched after time t to form (one of) the step-t edge(s), is the first element of F_t. g_t can be in any of the three graphs.
- $A_t(\boldsymbol{q})$ is the set of unmatched (hence "available") half-edges in G_q, the moment before g_t is matched, $\boldsymbol{q} \in \{\boldsymbol{x}, \boldsymbol{uv}, \boldsymbol{y}\}$. So $g_t \in A_t(\boldsymbol{q})$ for some $\boldsymbol{q}$.
- "Good" sets $K_t(\boldsymbol{z}), \boldsymbol{z} \in \{\boldsymbol{x}, \boldsymbol{u}, \boldsymbol{v}, \boldsymbol{y}\}$. Each of these contains both edges (that is, matched pairs of half-edges) and unmatched half-edges in F_t. $K_t(\boldsymbol{x}), K_t(\boldsymbol{u})$ denote the edges in the explored parts of $B_R(\boldsymbol{x})$ and $B_R(\boldsymbol{u})$ that correspond to each other, combined with unmatched half-edges in the frontier that have the potential to correspond once matched. In the same vein, $K_t(\boldsymbol{v}), K_t(\boldsymbol{y})$ denote the part where the two neighborhoods agree, and frontier half-edges that might agree. $K_t(\boldsymbol{x})$ are $K_t(\boldsymbol{u})$ are made to be isomorphic throughout time, so are $K_t(\boldsymbol{y})$ and $K_t(\boldsymbol{v})$.
- ι_t is an involution which restricts to isomorphisms $K_t(\boldsymbol{u}) \leftrightarrow K_t(\boldsymbol{x})$ and $K_t(\boldsymbol{v}) \leftrightarrow K_t(\boldsymbol{y})$ and maps corresponding edges and half-edges. We say ι_t maps a (half-)edge to its *partner* in the other graph.

Ideally, $K_t(\boldsymbol{x}) \cong K_t(\boldsymbol{u})$ is large, so the directed neighborhoods of $\boldsymbol{u}$ and $\boldsymbol{x}$ can be similar. Note that $K_t(\boldsymbol{z})$ are *merely labels on (half-)edges*. They don't affect the BFS probabilities, or how we add/remove edges from F_t or $A_t(\boldsymbol{q})$.

As we will see, the coupling results in $K_t(\boldsymbol{u}) \cap K_t(\boldsymbol{v}) = \varnothing$. This is needed because an (half-)edge in $K_t(\boldsymbol{u})$ imitates the $\boldsymbol{x}$-exploration. We don't want it to come under the influence of $\boldsymbol{v}$, which imitates $\boldsymbol{y}$, otherwise $\boldsymbol{x}$ and $\boldsymbol{y}$ will not be independent. Moreover, once an edge has been added to $K_t(\boldsymbol{z})$, it stay in all future $K_{t'}(\boldsymbol{z})$. A half-edge added to $K_t(\boldsymbol{z})$, however, can be removed at a later time.

Define $\beta = \bigcup_{z \in \{x,u,v,y\}} (\delta z \backslash \boldsymbol{z})$. At time t, an unmatched half-edge h_t that lives in the same graph as $\boldsymbol{z}$ is either

(1) in $K_t(\boldsymbol{z})$. We call h_t a *$\boldsymbol{z}$-good* half-edge.
(2) in $(F_t \cup \beta) \backslash K_t(\boldsymbol{z})$. Call h_t *$\boldsymbol{z}$-bad*.
(3) not in $F_t \cup \beta$ (so h_t attaches to a fresh vertex not yet discovered). Call h_t *fresh*.

Note that g_t can be both $\boldsymbol{u}$-good and $\boldsymbol{v}$-bad. Moreover, g_t cannot be both $\boldsymbol{u}$-good and $\boldsymbol{v}$-good.

5.3. **The coupling in real time: step by step.** We start the coupled exploration with $F_0 = \{\boldsymbol{x}, \boldsymbol{y}, \boldsymbol{u}, \boldsymbol{v}\}$ in sequence. $A_0(\boldsymbol{q})$ is the set of all half-edges in G_q. And $K_0(\boldsymbol{z}) = \boldsymbol{z}$. We will couple in such a way that for any $0 \leq t \leq t(R)$, g_t, the front of queue F_t, is never in $K_t(\boldsymbol{u}) \cup K_t(\boldsymbol{v})$. This is achieved by always placing the $K_t(\boldsymbol{x})$-half-edge h in front of its corresponding $K_t(\boldsymbol{u})$-half-edge $\iota(h)$ in F_t. And when we pop h from F_t, we delete $\iota(h)$ as well.

The evolution of F_t and $A_t(\boldsymbol{q})$ follows the standard BFS in G_q. If the an edge is formed by matching half-edges $j, k \in G_q$, $A_{t+1}(\boldsymbol{q}) = A_t(\boldsymbol{q}) \backslash \{j, k\}$. The other $A_t(\boldsymbol{q'})$ remain unchanged. We delete j, k from F_t (it's possible that one of them is not in F_t). And if k attaches to a fresh vertex w that have not been explore before, we append $\delta w \backslash \{k\}$ to the end of F_t. This is not affected by what is in $K_t(\boldsymbol{z})$.

The evolution of $K_t(\boldsymbol{z})$, however, requires some case-work and dictates how we couple. We have the following two types of g_t.

Case 1: $g_t \in K_t(\boldsymbol{x}) \cup K_t(\boldsymbol{y})$.

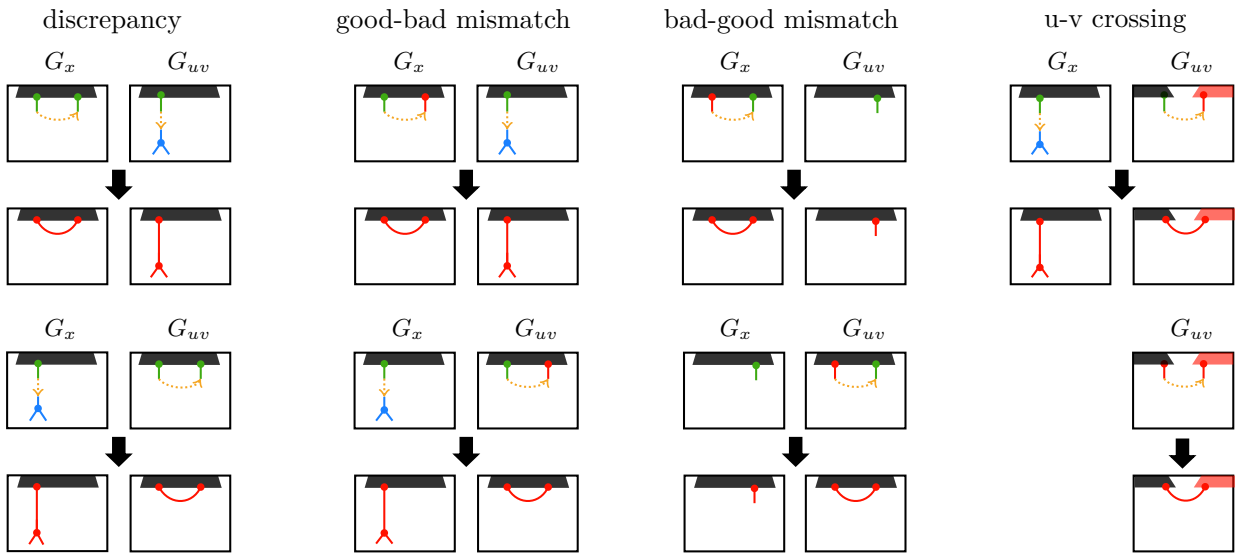

FIGURE 6. The result of each coupling error; black represents a good frontier half-edge, red represents a bad frontier half-edge, grey represents a fresh half-edge. The arrow indicates going from time t to $t+1$.

We will only detail the case $g_t \in K_t(\boldsymbol{x})$. The $g_t \in K_t(\boldsymbol{y})$ case is similar. So g_t is good, $\iota(g_t) \in K_t(\boldsymbol{u})$, $\iota(g_t) \in F_t$. With step t we will match g_t and $\iota(g_t)$ to form two edges. We want the two new edges to correspond under ι, but this is not always possible.

g_t matches each available half-edge (except g_t itself) with equal probability, namely, $p_t(\boldsymbol{x}) \equiv 1/(|A_t(\boldsymbol{x})| - 1)$. There are three kinds of available half-edge $h_t \in A_t(\boldsymbol{x})$ that g_t can match with:

1) $h_t \in K_t(\boldsymbol{x})\backslash\{g_t\}$ ("$\boldsymbol{x}$-good")
2) $h_t \in (F_t \cup \beta)\backslash K_t(\boldsymbol{x})$ ("$\boldsymbol{x}$-bad")
3) $h_t \notin F_t \cup \beta$ ("fresh")

We couple the matching of g_t with the matching of $\iota(g_t)$ in G_{uv}. The probability of $\iota(g_t)$ matching with every available half-edge except itself is $p_t(\boldsymbol{uv}) \equiv 1/(|A_t(\boldsymbol{uv})| - 1)$. There are similarly three kinds of $i_t \in A_t(\boldsymbol{uv})$:

a) $i_t \in K_t(\boldsymbol{u})\backslash\{\iota(g_t)\}$ ("$\boldsymbol{u}$-good")
b) $i_t \in (F_t \cup \beta)\backslash K_t(\boldsymbol{u})$ ("$\boldsymbol{u}$-bad")
c) $i_t \notin F_t \cup \beta$ ("fresh")

We want $\boldsymbol{u}$-exploration to imitate $\boldsymbol{x}$ as closely as possible. Note that $\iota(K_t(\boldsymbol{x}) \cap A_t(\boldsymbol{x})) = K_t(\boldsymbol{u}) \cap A_t(\boldsymbol{uv})$. So for each of the $|K_t(\boldsymbol{x}) \cap A_t(\boldsymbol{x})| - 1 = |K_t(\boldsymbol{u}) \cap A_t(\boldsymbol{uv})| - 1$ $\boldsymbol{x}$-good half-edge h_t, we can couple the g_t, h_t matching with the $\iota(g_t), i_t = \iota(h_t)$ matching, occurring with probability $\min(p_t(\boldsymbol{x}), p_t(\boldsymbol{uv}))$. This is our success case (1a).

Remark 5.1. For $i \in \{1,2,3\}$ and $k \in \{a,b,c\}$, we label a coupling case by (ik) if h_t is type i and i_t is type k. For $A, B, C, D \in \{\text{good}, \text{bad}, \text{fresh}\}$, we use $(A-B, C-D)$ to mean a A half-edge matching a B half-edge in G_x, while a C half-edge matching a D half-edge in G_{uv}.

(1a) *(good-good, good-good)*: A collision edge is created in the $\boldsymbol{x}$ neighborhood, and a corresponding collision edge in the $\boldsymbol{u}$ neighborhood.

$$K_{t+1}(\boldsymbol{x}) = (K_t(\boldsymbol{x})\backslash\{g_t, h_t\}) \cup (e_t : g_t \to h_t)$$
$$K_{t+1}(\boldsymbol{u}) = (K_t(\boldsymbol{u})\backslash\{\iota(g_t), \iota(h_t)\}) \cup (\iota(e_t) : \iota(g_t) \to \iota(h_t))$$

So $K_{t+1}(\boldsymbol{x})$ and $K_{t+1}(\boldsymbol{u})$ remain isomorphic.
(1a) occurs with probability $p_{t,1} \equiv (|K_t(\boldsymbol{x}) \cap A_t(\boldsymbol{x})| - 1) \cdot \min(p_t(\boldsymbol{x}), p_t(\boldsymbol{uv}))$.

At the same time, for each $\boldsymbol{x}$-good half-edge h_t, there is $\max(p_t(\boldsymbol{x}), p_t(\boldsymbol{uv})) - \min(p_t(\boldsymbol{x}), p_t(\boldsymbol{uv}))$ probability that correspondence fails due to probability discrepancy. This is *coupling error* case (1c/3a):

Remark 5.2. We construct the coupling such that whenever in one graph, the step is a collision that cannot be mirrored, the other graph will always find a fresh half-edge.

(1c/3a) *discrepancy: {(good-good, good-fresh), (good-fresh, good-good)}*: If $p_t(\boldsymbol{x}) \geq p_t(\boldsymbol{uv})$,[2] couple the collision edge $g_t \to h_t$ with $\iota(g_t) \to i_t$, where i_t is uniformly chosen from all fresh half-edges. If $p_t(\boldsymbol{x}) < p_t(\boldsymbol{uv})$, couple the collision edge $\iota(g_t) \to \iota(h_t)$ with $g_t \to j_t$ where j_t is fresh. In both cases, the two newly formed edges in G_x and G_{uv} do not agree. Good half-edges are corrupted. Nothing good is created.

$$K_{t+1}(\boldsymbol{x}) = K_t(\boldsymbol{x})\backslash\{g_t, h_t\}$$
$$K_{t+1}(\boldsymbol{u}) = K_t(\boldsymbol{x})\backslash\{\iota(g_t), \iota(h_t)\}$$

$\iota(h_t)$ remains in F_{t+1} but is removed from K_{t+1}: it becomes a "bad" half-edge since time $t+1$. If $p_t(\boldsymbol{x}) \geq p_t(\boldsymbol{uv})$ let i_t attach to fresh vertex w_u. The unmatched half-edges of w_u are added to F_{t+1}, but not $K_{t+1}(\boldsymbol{u})$: $\boldsymbol{u}$-bad. If $p_t(\boldsymbol{x}) < p_t(\boldsymbol{uv})$, let i_t attach to fresh vertex w_x. $\delta w_x\backslash j_t$ are $\boldsymbol{x}$-bad.
(1c/3a) occurs with probability
$p_{t,2} \equiv (|K_t(\boldsymbol{x}) \cap A_t(\boldsymbol{x})| - 1) \cdot (\max(p_t(\boldsymbol{x}), p_t(\boldsymbol{uv})) - \min(p_t(\boldsymbol{x}), p_t(\boldsymbol{uv})))$.

We have taken care of cases (1) and (a). It is also possible for g_t to match a $\boldsymbol{x}$-bad half-edge h_t that does not have a partner in G_{uv}. Since h_t is a frontier edge or incident to x, the resulting edge $g_t \to h_t$ must be a collision. $g_t \to h_t$ doesn't even have the potential to find a partner ($h_t \notin K_t(\boldsymbol{x})$), and will not be added to $K_{t+1}(\boldsymbol{x})$. This is *coupling error* case (2c):

(2c) *$\boldsymbol{x}$-good-bad mismatch: (good-bad, good-fresh)*: h_t is $\boldsymbol{x}$-bad. We couple collision $g_t \to h_t$ with $\iota(g_t) \to i_t$ where i_t is chosen uniformly from all fresh half-edges. Let i_t attach to vertex w_u. $\delta w_u\backslash i_t$ become $\boldsymbol{u}$-bad.

$$K_{t+1}(\boldsymbol{x}) = K_t(\boldsymbol{x})\backslash\{g_t\}$$
$$K_{t+1}(\boldsymbol{u}) = K_t(\boldsymbol{u})\backslash\{\iota(g_t)\}$$

(2c) occurs with probability $p_{t,3} = |((F_t \cup \beta)\backslash K_t(\boldsymbol{x})) \cap A_t(\boldsymbol{x})| \cdot p_t(\boldsymbol{x})$.

Similarly, $\iota(g_t)$ can hit a $\boldsymbol{u}$-bad half-edge i_t, and form an edge without a partner. This is *coupling error* case (3b):

(3b) *$\boldsymbol{u}$-good-bad mismatch: (good-fresh, good-bad)*: i_t is $\boldsymbol{u}$-bad. We couple collision $\iota(g_t) \to i_t$ with $g_t \to h_t$ where h_t is chosen uniformly from fresh half-edges and attached to vertex w_x. $\delta w_x\backslash h_t$ become $\boldsymbol{x}$-bad.

$$K_{t+1}(\boldsymbol{x}) = K_t(\boldsymbol{x})\backslash\{g_t\}$$
$$K_{t+1}(\boldsymbol{u}) = K_t(\boldsymbol{u})\backslash\{\iota(g_t)\}$$

(3b) occurs with probability $p_{t,4} \equiv |((F_t \cup \beta)\backslash K_t(\boldsymbol{u})) \cap A_t(\boldsymbol{uv})| \cdot p_t(\boldsymbol{uv})$.[3]

The remaining cases are (3) and (c): fresh half-edges. And the remaining probability is $1 - (p_{t,1} + p_{t,2} + p_{t,3} + p_{t,4})$ in both explorations. Even if $p_t(\boldsymbol{x})$ can be different from $p_t(\boldsymbol{uv})$, it does not matter which fresh half-edge we match because the result is isomorphic. So we have success case (3c):

(3c) *(good-fresh, good-fresh)*: h_t and i_t are both fresh vertices, attached to vertices w_x, w_u respectively.

$$K_{t+1}(\boldsymbol{x}) = (K_t(\boldsymbol{x})\backslash\{g_t, h_t\}) \cup (e_t : g_t \to h_t) \cup (\delta w_x\backslash h_t)$$
$$K_{t+1}(\boldsymbol{u}) = (K_t(\boldsymbol{u})\backslash\{\iota(g_t), \iota(h_t)\}) \cup (f_t : \iota(g_t) \to i_t) \cup (\delta w_u\backslash i_t)$$
$$F_{t+1} = (F_t\backslash\{g_t, \iota(g_t), h_t, i_t\}) \cup (\delta w_x\backslash h_t) \cup (\delta w_u\backslash i_t)$$

$\iota_{t+1}(h_t) = i_t,\ \iota_{t+1}(e_t) = f_t,\ \iota_{t+1}(\delta w_x\backslash h_t) = \delta w_u\backslash i_t$. Note that $\delta w_x\backslash h_t$ precedes $\delta w_u\backslash i_t$ in F_{t+1}, so a $K_{t+1}(\boldsymbol{u}) \cup K_{t+1}(\boldsymbol{v})$ half-edge will not appear in the front of the queue.
(3c) appears with probability $1 - (p_{t,1} + p_{t,2} + p_{t,3} + p_{t,4})$.

As we update $\boldsymbol{x}$ and $\boldsymbol{u}$, nothing happens in G_y. $A_t(\boldsymbol{y})$, $K_t(\boldsymbol{y})$, $K_t(\boldsymbol{v})$ do not change from time t to $t+1$.

[2]If $p_t(\boldsymbol{x}) = p_t(\boldsymbol{uv})$, (1a) will be mirrored perfectly, and discrepancy does not occur.

[3]It seems that we are wasting the fresh-half-edge probability by matching errors with a fresh half-edge. A more economical accounting is to combine (b2), (b3) and consider coupling (good-bad, good-bad). However, our excesses avoid additional casework. Our coupling is well-defined because there more than enough fresh half-edges at large: $p_{t,1} + p_{t,2} + p_{t,3} + p_{t,4} \leq \frac{|F_t \cup \beta|}{|A_t(\boldsymbol{x})|-1} + \frac{|F_t \cup \beta|}{|A_t(\boldsymbol{uv})|-1} \leq \frac{8+8\sqrt{n}\log n}{n-24\sqrt{n}\log n} \ll 1$ for any $0 \leq t \leq t(R)$. The middle inequality follows from $R \leq R_{\max}$.

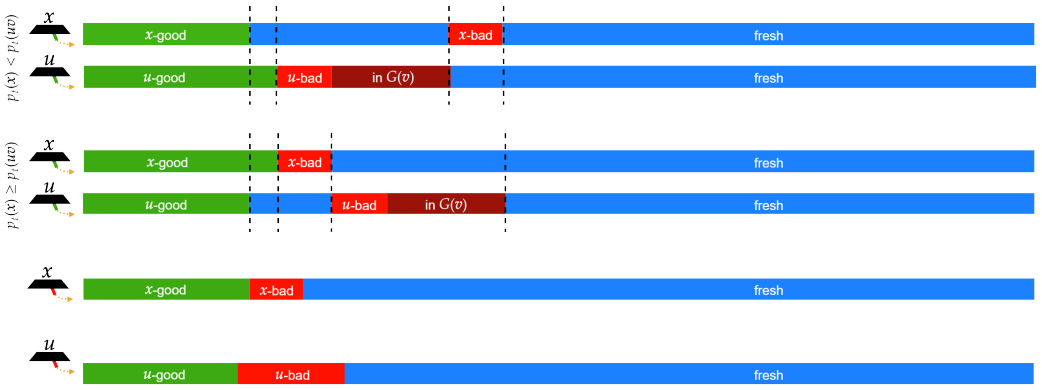

FIGURE 7. The distribution of h_t and $\iota(h_t)$ (if it exists) for a given g_t (to the left, good g_t is green, bad g_t is red). The length of a colored section is proportional to the numbers of available half-edges of such type. When g_t is good, we sample (h_t, i_t) according to the top two bars if $p_t(\boldsymbol{x}) > p_t(\boldsymbol{uv})$ and the 3rd and 4th otherwise. "in $G(\boldsymbol{v})$" is a prevalent subset of $\boldsymbol{u}$-bad.

Case 2: $g_t \notin K_t(\boldsymbol{x}) \cup K_t(\boldsymbol{y}) \cup K_t(\boldsymbol{u}) \cup K_t(\boldsymbol{v})$.

Here g_t can live in any of the three G_q. The edge formed by matching g_t to a uniformly chosen h_t has no potential for correspondence, or being included in any $K_t(\boldsymbol{z})$. Assume for concreteness $g_t \in G_x$ (the other two cases are similar). We only form one edge $g_t \to h_t$ at step t since g_t has no partner. h_t is uniformly chosen from any available half-edge in $A_t(\boldsymbol{x})$, good or bad or fresh.

If $h_t \in K_t(\boldsymbol{x})$ *(bad-good mismatch)*,

$$K_{t+1}(\boldsymbol{x}) = K_t(\boldsymbol{x}) \backslash \{h_t\}$$
$$K_{t+1}(\boldsymbol{u}) = K_t(\boldsymbol{u}) \backslash \{\iota(h_t)\}$$

Otherwise, $K_t(\boldsymbol{x})$ and $K_t(\boldsymbol{u})$ do not change from time t to $t+1$.

6. Distinguishability: analysis of the coupling, an introduction

Remind that we want to prove $\mathcal{E}_0(\boldsymbol{u}, \boldsymbol{v})$ does not hold with probability $o(n^{-2})$, for directions $\boldsymbol{u}, \boldsymbol{v}$ which we couple with $\boldsymbol{x}, \boldsymbol{y}$. In the BFS coupling, we gradually build the four directed neighborhoods altogether layer by layer: we find the depth-1 layers of *all four* neighborhoods, then the second layers of all four, then the third layers... Hence, after finding the $L_\circ$ layers of all four at step $t_\circ$, we can check if case (a) in $\mathcal{E}_0(\boldsymbol{u}, \boldsymbol{v})$ is satisfied. If not, $\mathcal{E}_0(\boldsymbol{u}, \boldsymbol{v})$ already holds, and we can stop the BFS exploration.

If (a) is satisfied by the top $L_\circ$ layers, that is, $B_{L_\circ}(\boldsymbol{u}) \cap B_{L_\circ}(\boldsymbol{u}) = \varnothing$ and both depth-$L_\circ$ directed neighborhoods are trees, we continue the BFS. We let $\mathcal{E}_1$ be the event that (a) holds.[4] As we make the following analysis, all based on BFS happenings after step $t_\circ$, *we will use property (a) as a given*, that is, we will assume event $\mathcal{E}_1$.

We call the probability in the $(\boldsymbol{x}, (\boldsymbol{u}, \boldsymbol{v}), \boldsymbol{y})$ joint distribution $\mathbb{P}_{\text{joint}}$, the probability in the $\boldsymbol{s}$-marginal distribution $P_{\boldsymbol{s}}$, where $\boldsymbol{s}$ is a subset of $\{\boldsymbol{x}, (\boldsymbol{u}, \boldsymbol{v}), \boldsymbol{y}\}$. If an event $\mathcal{E}'$ only concerns the graphs in $\boldsymbol{s}$, then $P_{\text{joint}}(\mathcal{E}') = P_{\boldsymbol{s}}(\mathcal{E}')$. For example, $P_{\boldsymbol{x},(\boldsymbol{u},\boldsymbol{v})}(\mathcal{E}_1 \wedge dist(B_R(\boldsymbol{x}), B_R(\boldsymbol{u})) \leq C) = P_{\text{joint}}(\mathcal{E}_1 \wedge dist(B_R(\boldsymbol{x}), B_R(\boldsymbol{u})) \leq C)$. Let $\mathbb{Q}(\mathcal{E}') \equiv \mathbb{P}_{\text{joint}}(\mathcal{E}_1 \wedge \mathcal{E}')$. For the rest of the analysis, we will be concerned with $\mathbb{Q}$, as only the case $\mathcal{E}_1$ requires further attention and BFS exploration.

Our goal is to prove that $\mathbb{Q}(\mathcal{E} \text{ and } B_R(\boldsymbol{u}) \cong B_R(\boldsymbol{u})) = \mathbb{P}_{\text{joint}}(\mathcal{E}, \mathcal{E}_1 \text{ and } B_R(\boldsymbol{u}) \cong B_R(\boldsymbol{u}))$
$= \mathbb{P}_{\boldsymbol{uv}}(\mathcal{E}, \mathcal{E}_1 \text{ and } B_R(\boldsymbol{u}) \cong B_R(\boldsymbol{u})) = \mathbb{P}(\mathcal{E} \text{ and not } \mathcal{E}_0(\boldsymbol{u}, \boldsymbol{v})) = o(n^{-2})$.

6.1. $\mathcal{E}_1$ implies that the depth-$L_\circ$ neighborhoods of $\boldsymbol{x}$ and $\boldsymbol{y}$ are both trees.

We prove, by induction on t, that in each instance in the joint distribution where $\mathcal{E}_1$ is satisfied, $B_{L_\circ}(\boldsymbol{x}) \cong B_{L_\circ}(\boldsymbol{y})$ are both trees, and the four depth-$L_\circ$ neighborhoods coincide with their respective $K_t(\boldsymbol{z})$ (they are all good).

[4] We drop the $(\boldsymbol{u}, \boldsymbol{v})$ from $\mathcal{E}_1(\boldsymbol{u}, \boldsymbol{v})$ since we only focus on the present $\boldsymbol{u}, \boldsymbol{v}$ that we couple.

$\mathcal{E}_1$ implies that in steps $0 \le t \le t_\circ$, the $\boldsymbol{u}$-exploration finds a fresh vertex or does nothing, and so does the $\boldsymbol{v}$-exploration. Assume (as induction hypothesis) that at time t before the t-th step, only success case (3c) has occurred before. This is true for $t = 0$. Then the explored parts of $\boldsymbol{x}$ and $\boldsymbol{u}$ are isomorphic trees with e_1 edges. The number of available (unmatched) half-edges in G_x is $nd - 2e_1$, strictly greater than the number of available half-edges in G_{uv}, as the $\boldsymbol{v}$-exploration also consumes half-edges. So $p_t(\boldsymbol{x}) < p_t(\boldsymbol{uv})$, and (good-good, good-fresh) will not happen. Furthermore, there is no bad edge in G_x. Referring back to top bar in Figure 7, if $\iota(g_t)$ in G_{uv} matches a fresh half-edge, g_t in G_x also matches a fresh half-edge: only success (3c) is possible. By induction, we have proved that if $\mathcal{E}_1$ holds, at every step $0 \le t \le t_\circ$ two new fresh vertices are found.

6.2. **The prevalence of error case (3b): some heuristics.** We want the coupled $B_R(\boldsymbol{x})$ and $B_R(\boldsymbol{u})$ to be similar, that is, the shared part $K_{t(R)}(\boldsymbol{x}) \cong K_{t(R)}(\boldsymbol{u})$ to be large. This happens when the number of coupling errors is small. Indeed, we will show that w.h.p., discrepancy (1c/3a), (2c) both happen $O_n(1)$ times.

However, the frequencies of a $\boldsymbol{u}$-good half-edge matching a $\boldsymbol{u}$-bad half-edge is much higher. If a $\boldsymbol{u}$-good half-edge $\iota(g_t)$ matches a half-edge i_t in the v-exploration (attached to vertex w_v), then the neighborhood of w_v, explored by $\boldsymbol{v}$, gets absorbed into $B_R(\boldsymbol{u})$ as well. Moreover, the edge $\iota(g_t) \to i_t$ and the w_t-neighborhood cannot be mirrored in the x-exploration – they are controlled by $\boldsymbol{y}$. Since the number of half-edges explored by $\boldsymbol{v}$ is similar to those by $\boldsymbol{u}$, we would expect such "crossing" $\iota(g_t) \to i_t$ happens as often as collisions within $B_R(\boldsymbol{u})$ take place: the frequency is a polynomial of $\log n$ in expectation. Unfortunately, there can be $B_R(\boldsymbol{u})$ cycles in the w_t-neighborhoods, making $B_R(\boldsymbol{x})$ and $B_R(\boldsymbol{z})$ further apart in cycle structure distance.

To make the notions of "crossing", "explored by $\boldsymbol{u}$", "$\boldsymbol{u}$-part" more precise, we introduce the following definitions.

Definition 6.1. We partition the vertices of $B_R(\boldsymbol{u}) \cup B_R(\boldsymbol{v})$ into 2 kinds: *u-vertices* $V(u)$ and *v-vertices* $V(v)$. We define them by the following: $u \in V(u)$, $v \in V(v)$. Every other vertex w in $B_R(\boldsymbol{u}) \cup B_R(\boldsymbol{v})$ is first discovered as a fresh vertex. If w is first discovered by a u-vertex, w a u-vertex, i.e. $w \in V(u)$. If w is first discovered by a v-vertex, w a v-vertex, i.e. $w \in V(v)$.

We let $G(\boldsymbol{u})$ be the induced subgraph of $B_R(\boldsymbol{u}) \cup B_R(\boldsymbol{v})$ on vertices $V(u)$, $G(\boldsymbol{v})$ be the induced subgraph of $B_R(\boldsymbol{u}) \cup B_R(\boldsymbol{v})$ on vertices $V(v)$

A *crossing* is an edge in $B_R(\boldsymbol{u}) \cup B_R(\boldsymbol{v})$ that is between a u-vertex and a v-vertex.

Remark 6.2. $G(\boldsymbol{u})$ and $G(\boldsymbol{v})$ are both connected. $B_R(\boldsymbol{u})$ can be partitioned into $G(\boldsymbol{u}), G(\boldsymbol{v})$, and crossings. A crossing can go in both directions: both a $V(u)$ half-edge matching a $V(v)$ half-edge, and a $V(v)$ half-edge matching a $V(v)$ half-edge. Both directions have the same effect: absorbing a neighborhood of $G(\boldsymbol{v})$ into $B_R(\boldsymbol{u})$.

A half-edge is added to $K_t(\boldsymbol{u})$ only in success case (3c), where the newly discovered vertex must be a u-vertex. An edge is added to $K_t(\boldsymbol{u})$ only if both half-edges are in $K_t(\boldsymbol{u})$. Thus, all the half-edges in $K_t(\boldsymbol{u})$ must attach to a u-vertex, and all the edges in $K_t(\boldsymbol{u})$ must be contained in $G(\boldsymbol{u})$. Any crossing or $G(\boldsymbol{v})$ edge cannot be in $K_{t(R)}(\boldsymbol{u})$.

Since $V(u) \cap V(v) = \varnothing, G(\boldsymbol{u}) \cap G(\boldsymbol{v}) = \varnothing$, we have proved that $K_t(\boldsymbol{u}) \cap K_t(\boldsymbol{v}) = \varnothing$ for every $0 \le t \le t(R)$. A discovered u-vertex can have all, some, or none of its attached half-edges in $K_t(\boldsymbol{u})$.

6.3. **Bounding $dist(\mathcal{C}(B_R(\boldsymbol{x})), \mathcal{C}(B_R(\boldsymbol{u})))$ by the number of edge cuts: a prelude.**

Lemma 6.3. *G_1, G_2 are two rooted graphs. If we can delete k edges from G_1, such that the resulting graph (may be disconnected) is a G_3 containing the root and a forest (that is, parts not containing the root are all trees), where G_3 has the same cycle structure as G_2, then $dist(\mathcal{C}(G_1), \mathcal{C}(G_2)) \le k$.*

Proof. Let S be the set of G_1 edges deleted. Start with G_1, we recursively delete edges that are both in S and in a cycle in the remaining graph. If there is a cycle in what remains but not in G_3, there must be an edge in S that can destroy the cycle. Delete this edge, which does not disconnect the graph as the edge is in the cycle. So by deleting $k' \le l$ edges from G_1 we get a connected graph G_3' with the same cycle structure as G_3, and hence as G_2. It requires no more than k' delete operations in $\mathcal{C}(G_1)$ to remove these k' edges. So $dist(\mathcal{C}(G_1), \mathcal{C}(G_2)) \le k' \le k$. □

Let $z \in \{x, u, v, y\}$. Assume that at the end of time, deleting $k(\boldsymbol{z})$ edges from $B_R(\boldsymbol{z})$ will result in a graph that is a forest combined with a subgraph containing vertex z, where the latter has the same cycle structure as $K_{t(R)}(\boldsymbol{z})$. Then $dist(\mathcal{C}(B_R(\boldsymbol{z})), \mathcal{C}(K_{t(R)}(\boldsymbol{z})) \leq k(\boldsymbol{z})$. Therefore, in every instance of the joint distribution

$$dist(\mathcal{C}(B_R(\boldsymbol{x})), \mathcal{C}(B_R(\boldsymbol{u}))) \leq k(\boldsymbol{u}) + k(\boldsymbol{x})$$
$$dist(\mathcal{C}(B_R(\boldsymbol{y})), \mathcal{C}(B_R(\boldsymbol{v}))) \leq k(\boldsymbol{v}) + k(\boldsymbol{y})$$

In our proof, will only bound $k(\boldsymbol{u}), k(\boldsymbol{x})$ as $k(\boldsymbol{v}), k(\boldsymbol{y})$ are similar. The goal is to prove the rather surprising result: $k(\boldsymbol{u}), k(\boldsymbol{x}) \leq O_n(1)$ with probability $1 - o(n^{-2})$ in the joint distribution. In other words, in most (a $1 - o(n^{-2})$ portion) of the instances in the joint distribution, our coupling is very successful.

Of course, we will choose carefully which edges in $B_R(\boldsymbol{u})$ and $B_R(\boldsymbol{x})$ to cut. We will start with reducing the $G(\boldsymbol{v})$ part of $B_R(\boldsymbol{u})$ into a forest, that is, effectively erasing the $G(\boldsymbol{v})$ part from $B_R(\boldsymbol{u})$'s cycle structure. Armed with Lemma 6.3, we do not have to worry about disconnecting the graph as we cut edges.

7. Coupling analysis: eliminate $G(\boldsymbol{v})$ from $\mathcal{C}(B_R(\boldsymbol{u}))$ by cutting edges

Let $E^{\times}_{G(\boldsymbol{v})}$ be the number of edges we cut to eliminate $G(\boldsymbol{v})$ from $B_R(\boldsymbol{u})$'s cycle structure. In this section we prove the following proposition,

Proposition 7.1.

$$\mathbb{Q}(E^{\times}_{G(\boldsymbol{v})} > C_2 \equiv 66 + 33 \cdot 630 \cdot (33 + 2^{22})^2) = o(n^{-2})$$

We will first define crossing neighborhoods, "balls" in $G(\boldsymbol{v})$ that a crossing brings into $B_R(\boldsymbol{v})$, and establish that $B_R(\boldsymbol{v}) \cap G(\boldsymbol{v})$ is actually the union of these "balls".

Then we will first remove the cycles completely contained within one neighborhood by removing short cycles (see lemma 2.14). Then we will disconnect, or sever, each pair of intersecting crossing neighborhoods. This removes all cycles of $B_R(\boldsymbol{u})$ containing a v-vertex.

Disconnection requires, w.h.p., only $O(1)$ edge cuts because crossing neighborhoods are shallow (with depth at most $\frac{7}{16}\log_{d-1} n$) and can be w.h.p. disconnected by $O(1)$ edges. Another important reason is that w.h.p. the number of intersecting pairs is $O(1)$. We prove this by differentiating upper and lower crossings, and examine the step-by-step BFS probabilities.

7.1. The part of $B_R(\boldsymbol{u})$ that is explored by $\boldsymbol{v}$, at the end of time.

Let H be the end-of-time induced subgraph of $B_R(\boldsymbol{u})$ on $V(v)$, the subgraph in question. H does not contain crossings. Edges in H are all $\boldsymbol{u}$-bad. We introduce the definitions to describe $H = B_R(\boldsymbol{u}) \cap G(\boldsymbol{v})$.

Definition 7.2. Suppose we are at the end of time. Let e be a crossing, with its $\boldsymbol{u}$-vertex s_e and $\boldsymbol{v}$-vertex t_e. Define the *crossing neighborhood of* e, $N(e)$, to be the subgraph of $G(\boldsymbol{v})$ induced by $\{w \in V(v), dist_{G(\boldsymbol{v})}(w, t_e) \leq R - \text{depth}(s_e) - 1\}$, in which edges between distance $(R - \text{depth}(s_e) - 1)$-vertices are deleted. We rooted $N(e)$ at t_e. Here, $\text{depth}(s_e)$ refers to the depth in $G(\boldsymbol{u})$. $dist_{G(\boldsymbol{v})}$ refers to the length of the shortest path in $G(\boldsymbol{v})$, not in $B_R(\boldsymbol{u}) \cup B_R(\boldsymbol{v})$.

Remark 7.3. Note that the definition of $N(e)$ stay fixed: we denote the remaining graph after round k of edge cutting by $N(e) \cap H_k$, where we define H_k to be a "mask" lacking the edges we cut. This ensures that across different crossing neighborhoods, the *same* set of edges are cut. Without the notation, keeping track of what remains in each neighborhood will be difficult. And this isn't just notational convenience: if edge f can be cut in $N(e_1)$ but not $N(e_2)$, our proof could fail.

If edge f with vertices w_1, w_2 is in $H \subseteq B_R(\boldsymbol{u})$, then WLOG $dist_{B_R(\boldsymbol{u})}(w_1, u) \leq R - 1$. So there is a path $u \to w_1$ of length $\leq R - 1$, travelling through some crossing e between s_e and t_e. By definition of $N(e)$, $f \in N(e)$. Conversely, if $f \in N(e)$ for some $N(e)$, find the shortest path from f to t_e then through e to u, then $f \in B_R(\boldsymbol{u})$. This gives us the following proposition:

Lemma 7.4. *H, the induced subgraph of $B_R(\boldsymbol{u})$ on $V(\boldsymbol{v})$, is equal to $\cup_{crossing\ e} N(e)$*

We want to remove every cycle in $B_R(\boldsymbol{u})$ that includes a v-vertex, so the remaining cycle structure will be contained in $G(\boldsymbol{u})$. A cycle either includes v-vertices only (so it is in H), or both u- and v-vertices. We further divide the former into two kinds, resulting in three types of such cycles (see Figure 8):

(1) contained in H, completely contained in one $N(e)$

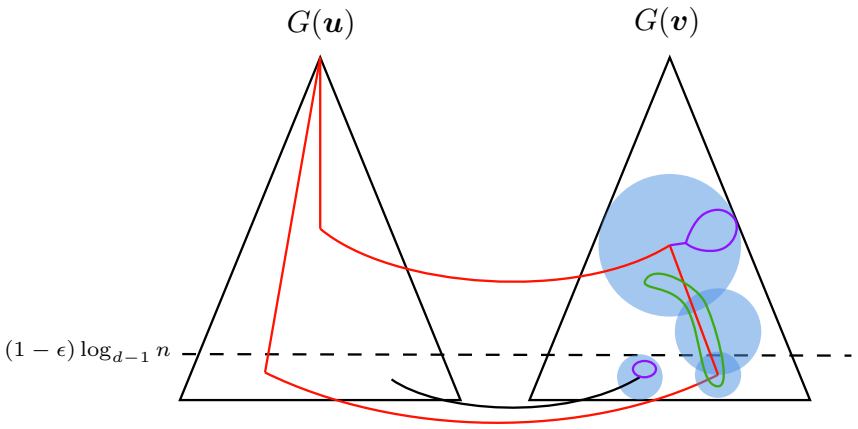

FIGURE 8. 3 types of cycles and how to bound the number of edge cuts that eliminate them:
purple: type (1); Lemma 2.14 (short cycle bound)
green: type (2); Lemma 2.13 (shallow collision bound), Definition 7.7
red: type (3); subsection 7.4 (counting the number intersecting crossing neighborhoods), Definition 7.10 (upper vs. lower crossing)

(2) contained in H, traversing several $N(e)$
(3) includes both $V(\boldsymbol{v})$ and $V(\boldsymbol{u})$ vertices

In the next 2 subsections, we will describe how we, *at the end of time*, will erase these cycles. In subsection 7.4, we count the number of edges we cut in the whole process, and show w.h.p. the number of cuts is $O_n(1)$. We start with cycle type (1).

7.2. **Round 1: remove within-neighborhood cycles by popping all short cycles.** Since each crossing neighborhood $N(e)$ has radius $\leq R - L_\circ - 1 \leq \frac{7.1}{16}\log_{d-1} n$, we might hope that all the cycles within $N(e)$ are *short* (cycle length $\leq 2R - 2L_\circ$). Lemma 3.11 tells us that with probability $1 - o(n^{-2})$, we can pop all the short cycles in G_{uv} by cutting no more than 66 edges.

However, $N(e)$ can contain longer cycles: imagine the roughly $n^{7/16}$ vertices with distance $\frac{7}{16}\log_{d-1} n$ from t_e form a loop. Fortunately, if an edge in $N(e)$ is part of any cycle, it is also be part of a short cycle.

Lemma 7.5. *Let N be a graph rooted in t where each vertex is of distance $\leq r$ from the root. Then if edge e is in some cycle C, there must be some cycle C' with length $\leq 2r+1$ such that $e \in C'$.*

The proof is in Appendix A. We consider $\mathcal{C}_0(\boldsymbol{v})$, the cycle structure of $G(\boldsymbol{v})$ with length $\leq 2R - 2L_0$, and do the first round of edge-cutting to reduce H to H_1. H_1 may be disconnected.

Corollary 7.6. *We cut $\gamma(\mathcal{C}_0(\boldsymbol{v}))$ edges such that there is no cycle, short or long, in what remains of $\mathcal{C}_0(\boldsymbol{v})$. Let the remaining part of H be H_1. Then, for any crossing e, there is no cycle, short or long, contained in $N(e) \cap H_1$ (that is, all edges in $N(e)$).*

Proof. $2(R - depth(s) - 1) + 1 < 2R - 2L_0$ for any crossing e. If there is a cycle $C \subset N(e) \cap H_1$ in the post-cutting graph, then C is in $N(e)$. By Lemma 6.15, every edge of C is in the pre-cutting $\mathcal{C}_0(\boldsymbol{v})$ as well. Since C survives the cutting intact, C is a cycle in the post-cutting $\mathcal{C}_0(\boldsymbol{v})$. However, we cut $\gamma(\mathcal{C}_0(\boldsymbol{v}))$ to reduce $\mathcal{C}_0(\boldsymbol{v}))$ to a tree. Contradiction. □

We have popped all type (1) cycle. If there is a cycle C contained in H_1, it cannot be completely contained in any single $N(e)$. Therefore, there exist two crossings $e \neq f$, such that C includes vertex $w_e \in N(e) \backslash N(f)$ and vertex $w_f \in N(e) \backslash N(f)$.

If there is a cycle $C' \subseteq G(\boldsymbol{u}) \cup \{\text{crossings}\} \cup H_1$, and C' includes both u- and v-vertices (i.e. type (3)), then C' contains two crossings e, f as edges, and a path in $G(\boldsymbol{v})$ that connects t_e, t_f (the path is unnecessary if $t_e = t_f$).

We will pop these two kinds of cycles in the next subsection, by severing every pair of intersecting neighborhoods.

7.3. **Round 2: sever crossing neighborhoods, and remove the rest of cycles that touches $G(\boldsymbol{v})$.**

Definition 7.7. $\mathcal{E}_4$ is the event, for a fixed d-regular graph G, that for every two vertices $t_1 \neq t_2$ of G, and every pair of radii $r_1, r_2 \in [0, R - L_\circ - 1]$, the following holds: we BFS explore the neighborhoods (all d-directions) of t_1, t_2 up to depth r_1, r_2 respectively, and define $V_{r_1}(t_1), V_{r_2}(t_2)$, and $G_{r_1}(t_1), G_{r_2}(t_2)$ (ref Definition 6.1);[5] the number of crossings between $G_{r_1}(t_1)$ and $G_{r_2}(t_2)$ is no more than 33.

Definition 7.8.

Let H be a subgraph of G (H could be disconnected). For vertices $t_1 \neq t_2$, consider subgraphs H_1 and H_2 of G, $t_1 \in H_1, t_2 \in H_2$ (both could be disconnected). If the following two properties are satisfied, we say that H_1 and H_2 are *severed by* H:

(1) There is no path in $(H_1 \cup H_2) \cap H$ between any vertices $w_1 \in H_1 \backslash H_2$ and $w_2 \in H_2 \backslash H_1$. w_1, w_2 do not have to be in H.
(2) There is no path in $(H_1 \cup H_2) \cap H$ between t_1, t_2.

Lemma 7.9. *At the end of time, assume that $\mathcal{E}_4$ holds on G_{uv}. For every two crossing neighborhoods $N(e)$ and $N(f)$ in $G(\boldsymbol{v})$, with $t_e \neq t_f$, we can delete no more than 33 edges from H_1 to result in an $H'(e, f)$ that severs $N(e), N(f)$.*

Proof. Let $r_e = R - depth(s_e) - 1$, $r_f = R - depth(s_f) - 1$, S be the set of crossing edges between $G_{r_e}(t_e)$ and $G_{r_f}(t_f)$ in G_{uv}. $|S| \leq 33$. Cut $S \cap H_1$ from H_1. If $w_1 \in N(e) \backslash N(f)$, then clearly $w_1 \notin V_{r_f}(t_f)$. Similarly $w_2 \notin V_{r_e}(t_e)$. Also, $t_e \in V_{r_f}(t_f)$ and $t_f \in V_{r_f}(t_f)$. In order for a point in $V_{r_e}(t_e)$ and point $V_{r_f}(t_f)$ to be connected, the path must pass through a crossing. Hence (1) (2) are satisfied. □

Given H_1, we will sever every pair of $N(e), N(f)$ with $t_e \neq t_f$ that intersect at at least one vertex. For $N(e), N(f)$ with $t_e = t_f$, we will simply remove crossing f. This is round 2 of our edge cuts, resulting in $H_2 \subset H_1$.

After cutting these edges, there will be no cycle left in $G(\boldsymbol{v}) \cap H_2$ (because neighborhoods are severed). There is also no cycle that includes both u- and v-vertices, as each $t_e \neq t_f$ are disconnected in H_2, and t_e can only connect to one crossing.

Note that if $\mathcal{E}_4$ holds, with each intersecting pair (we count $t_e = t_f$ pairs as intersecting) we delete no more than 33 edges. $\mathbb{Q}(\neg\mathcal{E}_4) = o(n^{-2})$ by Lemma 2.12 and a union bound over all u, v, r_1, r_2. We only need to prove that, independent of $\mathcal{E}_4$, the number of intersecting $N(e), N(f)$ pairs is $O(1)$ with probability $1 - o(n^{-2})$.

7.4. **Upper bounding the number of intersecting neighborhoods.** The goal of this section is to prove that at the end of time in the joint distribution, $\mathbb{Q}(|(e, f) : N(e) \cap N(f) \neq \varnothing| > 630 \cdot (33 + 2^{22})^2) = o(n^{-2})$.

To do so we differentiate between two types of crossing neighborhoods:

Definition 7.10. A crossing e is a *upper crossing* if $\text{depth}(s_e) \leq (1 - \epsilon) \log_{d-1} n$. A crossing e is a *lower* crossing if $\text{depth}(s_e) > (1 - \epsilon) \log_{d-1} n$. $\epsilon = 1/64$.

Remark 7.11. An upper crossing creates a large neighborhood in $G(\boldsymbol{v})$, with radius as much as $R - L_\circ - 1 < \frac{7.1}{16} \log_{d-1} n$, and these neighborhoods are more likely to intersect. However, there are not many (w.h.p. $O(1)$) of them because they exist between two shallow neighborhoods (see Lemma 2.13).

A lower crossing creates a small neighborhood with radius $\leq 1.1\epsilon \log_{d-1} n$, and these neighborhoods are sparsely positioned, rarely intersecting. However, there are many of them: at least $\log_{d-1} n$ in expectation, cf. Lemma 2.12.

Our previous analyses are all end-of-time. However, here we bound the number of pairs *at the end of time, by focusing on each BFS step* as the BFS evolves. And we will look at *all instances in the joint distribution* that satisfy $\mathcal{E}_1$. In each step, we have *no knowledge of the future*.[6] Time t is the moment we are about to work out the t-th step. At time t, we can only see what we have explored in the past $t - 1$ steps.

[5] $G_{r_1}(t_1), G_{r_2}(t_2)$ are induced subgraphs of $B_{r_1}(t_1) \cap B_{r_2}(t_2)$ on $V_{r_1}(t_1), V_{r_2}(t_2)$ respectively.

[6] We will make sure everything we define is measurable at time t. Conditioning on an end-of-time event (or equivalently, restricting our attention only to certain instances of the distribution) will mess up the probabilities at each prior BFS step. We avoid doing so by introducing indicator functions of each step t.

Definition 7.12. Let G_t be the explored part of G_{uv} at time t, determined by the first $t-1$ steps (it will eventually be a subgraph of $B_R(\boldsymbol{u}) \cup B_R(\boldsymbol{v})$). Let $V_t(\boldsymbol{v})$ be the set of v-vertices in G_t at time t, $G_t(\boldsymbol{v})$ be the induced subgraph of G_t on $V_t(\boldsymbol{v})$.

Each step also updates the distance between two vertices. As t increases and more edges are connected, the distance in $G_t(\boldsymbol{v})$ decreases or stays the same.

Definition 7.13. We define $dist_t(w_1, w_2)$, where $w_1, w_2 \in V_t(\boldsymbol{v})$, to be the length of the shortest path between w_1, w_2 in $G_t(\boldsymbol{v})$.[7] For a crossing e present at time t, define the crossing neighborhood at time t, $N_t(e)$, to be the subgraph of $G_t(\boldsymbol{v})$ induced by $\{w \in V_t(\boldsymbol{v}),\ dist_t(w, t_e) \le R - \text{depth}(s_e) - 1\}$, in which edges between distance $(R - \text{depth}(s_e) - 1)$-vertices are deleted. We root $N_t(e)$ in t_e.

Remark 7.14. Our previously defined $N(e) = N_{t(R)+1}(e)$. $N_t(e)$ is also defined by the distance at time t, not end of time distance. A vertex, present at time t but not in $N_t(e)$, could be added to $N_{t'}(e)$ at a later time. If $N_t(e) \cap N_t(f) \neq \varnothing$, then $N_{t'}(e) \cap N_{t'}(f) \neq \varnothing$ for all $t \le t' \le t(R) + 1$.

Definition 7.15. Let Y_t be the number of unordered (e, f) pairs where e, f are crossings present after step t, $N_{t+1}(e) \cap N_{t+1}(f) \neq \varnothing$, $N_t(e) \cap N_t(f) = \varnothing$ (the last condition is vacuously satisfied if e or f does not exist at time t).

Essentially, Y_t is the number of intersecting $(N(e), N(f))$ pair "created" by step t. The quantity we want to bound, the total number of intersecting $(N(e), N(f))$ pairs, is $\sum_{t_\circ \le t \le t(R)} Y_t$.

Definition 7.16. We keep track of these random variables in the BFS evolution. They are all $\mathcal{F}_t$-measurable, where $\mathcal{F}_t$ is the σ-algebra generated by the coupling up to time t.

- U_t: the number of upper crossings at time t (before the t-th step takes place)
- L_t: the number of lower crossings at time t
- X_t: at time t, $\max_{w \in V_t(\boldsymbol{v})} |\text{lower crossing } e : dist_t(t_e, w) \le 2.2\epsilon \log_{d-1} n|$, that is, the maximum number of lower crossings whose t_e are all close to a single vertex

Let the the $(t_1 - 1)$-th step complete the depth-$(1-\epsilon)\log_{d-1} n$ explorations on $\boldsymbol{u}, \boldsymbol{v}$. t_1 is a stopping time. The number of (upper-upper) pairs no longer change since time t_1, and the quantity is $\sum_{t_\circ \le t < t_1} Y_t \le U_{t_1}(U_{t_1}+1)/2$. As we consider lower crossings, we only need to consider the BFS at each step $t_1 \le t \le t(R)$.

Now we will analyze the BFS at each step t, $t_1 \le t \le t(R)$.Step t creates in G_{uv} either no edge, or one edge $e_t : g_t \to h_t$. In fact, the majority of steps t has $Y_t = 0$. If $Y_t > 0$, there are crossings e, f whose neighborhoods are made intersect by step t. An edge e_t is created in G_{t+1}. It must be one of the three cases:

(1) e_t is a u-to-v crossing, that is, g_t attaches to u-vertex s_{e_t}; h_t attaches to v-vertex t_{e_t}. e_t does not update distance as $G_t(\boldsymbol{v}) = G_{t+1}(\boldsymbol{v})$. No two existing crossing neighborhoods will be joined by e_t.
   - In order for case (1) to happen, g_t must attach to a u-vertex.
   - In order for $N_t(e_t)$ and $N_t(f)$ to intersect, where f is a previous crossing, h_t must satisfy $dist_t(t_{e_t}, t_f) \le 1.1\epsilon \log_{d-1} n + (R - \text{depth}(f) - 1)$.
   - At time t, the number of half-edges in the frontier that satisfy the h_t requirement is no more than $U_t \cdot n^{7.1/16 + 1.1\epsilon} + L_t \cdot n^{2.2\epsilon}$.
   - Here we upper bound Y_t, that is, the number of existing crossing neighborhoods $N_{t+1}(e_t)$ can intersect. The number of upper neighborhoods $N_{t+1}(e_t)$ intersects is $\le U_t$. If after step t, $N_{t+1}(e_t)$ intersects lower neighborhoods $N_{t+1}(f_1), \cdots, N_{t+1}(f_k)$, $k > 1$, then before step t (i.e. at time t), $dist_t(t_{e_t}, t_{f_i}) \le 2.2\epsilon \log_{d-1} R$ for every $1 \le i \le k$. By definition $k \le X_t$. Thus, in every instance of the joint distribution, $Y_t \le U_t + X_t$.

(2) e_t is a v-to-u crossing, that is, g_t attaches to v-vertex t_{e_t}; h_t attaches to u-vertex s_{e_t}. Similar to (1),
   - g_t must attach to a v-vertex. In order for $N_t(e_t)$ and $N_t(f)$ to intersect, where f is a previous crossing, $dist_t(t_{e_t}, t_f) \le 1.1\epsilon \log_{d-1} n + (R - \text{depth}(f) - 1)$.
   - h_t must attach to a u-vertex.
   - The number of half-edges at time t satisfying the h_t-requirement is no more than $n^{1/2} \log_{d-1} n$.
   - In every instance, $Y_t \le U_t + X_t$.

[7] Not length of the shortest path in G_t, i.e. the explored part of $B_R(\boldsymbol{u}) \cup B_R(\boldsymbol{v})$. This accords with our definition of end-of-time crossing neighborhoods, which relies on the end-of-time $G(\boldsymbol{v})$ distance, not $B_R(\boldsymbol{u}) \cap B_R(\boldsymbol{v})$ distance.

(3) e_t is not a crossing. No new crossing neighborhood appears, but two existing neighborhoods are joined together by e_t. Let e_t have vertices w_1, w_2. In order for $N_{t+1}(e), N_{t+1}(f)$ to be joined, we must have $w_1 \in N_t(e), w_2 \in N_t(f)$ (or we can flip the labels e, f).
- g_t must attach to a v-vertex, which is in at least one crossing neighborhood $N_t(e)$ at time t.
- h_t has the same requirement as g_t.
- The number of half-edges satisfying the h_t-requirement is no more than $U_t \cdot n^{7.1/16} + L_t \cdot n^{1.1\epsilon}$.
- Let all the intersecting pairs created by e_t be $(N_{t+1}(e_1), N_{t+1}(f_1)), \cdots, (N_{t+1}(e_k), N_{t+1}(f_k))$. Let $w_1 \in N_t(e_i)$, $w_2 \in N_t(f_i)$, $1 \le i \le k$ (or we flip e_i, f_i). Among $N_t(e_i)$, all the lower neighborhoods $N_t(e_j)$ satisfy $dist_t(t_{e_j}, w_1) \le 2.2\epsilon \log_{d-1} n$. Here the distance is that of time t. So among all e_i, there are at most $U_t + X_t$ *distinct* crossings. And there are at most $U_t + X_t$ *distinct* f_i. In every instance, $Y_t = k \le (U_t + X_t)^2$.

Let g_t be the half-edge to be matched at step t, if g_t, satisfies the g_t-requirement of case (i), we say it is *i-good*. If h_t, the half-edge g_t matches to at step t, satisfies the g_t requirement of case (i), we say it is *i-good*.

Instead of bounding $\sum_{t_1 \le t \le t(R)} Y_t$, we define Y_t's proxy Z_t, which is automatically 0 when U_t or X_t is large. So Z_t's sum can be stochastically dominated by binomials. Let $C_1 \equiv 2^{22}$.

$$\mathbb{1}(\mathcal{E}_t) \equiv \mathbb{1}(U_t \le 33) \cdot \mathbb{1}(L_t \le (4e \log_{d-1} n)^2) \cdot \mathbb{1}(X_t \le C_1)$$
$$Z_t \equiv Y_t \cdot \mathbb{1}(\mathcal{E}_t)$$

$$\begin{aligned}\sum_{t_1 \le t \le t(R)} Z_t &= \sum_{t_1 \le t \le t(R)} \mathbb{1}(Y_t > 0) \cdot Z_t \\ &= \sum_{i=1,2,3} \sum_{t_1 \le t \le t(R)} \mathbb{1}(e_t \text{ is case i}) \cdot Z_t \\ &\le (33 + C_1)^2 \cdot \sum_{i=1,2,3} \sum_{t_1 \le t \le t(R)} \mathbb{1}(g_t \text{ i-good}) \cdot \mathbb{1}(h_t \text{ i-good}) \cdot \mathbb{1}(\mathcal{E}_t)\end{aligned}$$

The last inequality follows from the fact that $Z_t \le (33 + C_1)^2$ for *any t in every instance of the joint distribution*, as a result of Z_t's definition and our case analysis. We will separately bound each of the $i = 1, 2, 3$ sums. *All* the following inequalities will hold for *any $t_1 \le t \le t(R)$ in every possible instance of the distribution (that is, they hold almost surely).*

Case 1: at each t, for the inequality below, we can plug in the bounds $U \le 33, L_t \le (4e \log_{d-1} n)^2$ in our upper bound for the number of 1-good h_t: if the bounds fail, $\mathbb{1}(h_t \text{ 1-good}) \cdot \mathbb{1}(\mathcal{E}_t)$ is automatically 0.

$$\mathbb{Q}(\mathbb{1}(h_t \text{ 1-good}) \cdot \mathbb{1}(\mathcal{E}_t) = 1) \le \frac{33 \cdot n^{7.1/16 + 1.1\epsilon} + (4e \log_{d-1} n)^2 \cdot n^{2.2\epsilon}}{nd - 2n^{1/2} \log_{d-1} n} \le n^{-0.51}$$

At the same time, $\sum_{t_1 \le t \le t(R)} \mathbb{1}(g_t \text{ 1-good}) \le n^{1/2} \log_{d-1} n$, since an 1-good g_t must attach to a u-vertex. Thus, the following random variable, which is distributed across all instances in the joint distribution satisfying $\mathcal{E}_1$,

$$\sum_{t_1 \le t \le t(R)} \mathbb{1}(g_t \text{ 1-good}) \cdot \mathbb{1}(h_t \text{ 1-good}) \cdot \mathbb{1}(\mathcal{E}_t) = \sum_{t_1 \le t \le t(R),\ g_t \text{1-good}} \mathbb{1}(h_t \text{ 1-good}) \cdot \mathbb{1}(\mathcal{E}_t)$$

is stochastically dominated by $A_1 \equiv \text{Bin}\left(n^{1/2} \log_{d-1} n, n^{-0.51}\right)$. [8]

[8]The stochastic dominance actually follows from the following coupling. Let $W_t = \mathbb{1}(h_t 1 - good) \cdot \mathbb{1}(\mathcal{E}_t)$ be a random variable. As the BFS evolves step by step, if $\mathbb{1}(g_t \text{ 1-good}) = 0$, we couple W_t with $\overline{W}_t = 0$ a.s.; otherwise we couple W_t with $\overline{W}_t = \text{Ber}(n^{-0.51})$. The latter is possible because $\mathbb{Q}(W_t = 1) \le n^{-0.51}$ *in every possible instance of the BFS evolution, no matter what happened in the first $t-1$ steps*, so the distribution of $\overline{W}_t$ can be *made independent of the first $t-1$ steps*. After BFS finishes, define $W_{t(R)+1} = 0$ a.s. and $\overline{W}_{t(R)+1}$ to be $\text{Ber}(n^{1/2} \log_{d-1} n - \sum_t \mathbb{1}(g_t \text{ 1-good}), n^{-0.51})$. $\overline{W}_{t(R)+1}$ is given only $\sum_t \mathbb{1}(g_t \text{ 1-good})$. Its distribution is independent of each Bernoulli $\overline{W}_t$ in the previous steps. As a result, for every $t_1 \le t \le t(R) + 1$, in every instance in the joint distribution of $(W_1, \cdots, W_t, \overline{W}_1, \cdots, \overline{W}_t)$, $\sum_{t_1 \le i \le t} W_i \le \sum_{t_1 \le i \le t} \overline{W}_i$. The marginal distribution of $\sum_{t_1 \le t \le t(R)+1} W_i$ is that of $\sum_{t_1 \le t \le t(R),\ g_t \text{1-good}} \mathbb{1}(h_t \text{ 1-good}) \cdot \mathbb{1}(\mathcal{E}_t)$ in the joint distribution of the $(\boldsymbol{x}, \boldsymbol{uv}, \boldsymbol{y})$ coupling. The marginal distribution of $\sum_{t_1 \le t \le t(R)+1}$, due to independence, is $\text{Bin}\left(n^{1/2} \log_{d-1} n, n^{-0.51}\right)$.

Case 2: a similar analysis gives

$$\mathbb{Q}(\mathbb{1}(h_t \text{ 2-good}) = 1) \le \frac{n^{1/2}\log_{d-1} n}{nd - 2n^{1/2}\log_{d-1} n} \le n^{-0.499}$$

$$\sum_{t_1 \le t \le t(R)} \mathbb{1}(g_t \text{ 2-good}) \cdot \mathbb{1}(\mathcal{E}_t) \le 33 \cdot n^{7.1/16+1.1\epsilon} + (4e\log_{d-1} n)^2 \cdot n^{2.2\epsilon} \le n^{0.49}$$

The display above results from our requirement for g_t, and the definition of $\mathcal{E}_t$: once U_t exceeds 33 or L_t exceeds $(4e\log_{d-1} n)^2$, any 2-good g_t produced by a new crossing will not be counted. Hence, only g_t produced by older crossings show up in this sum.

Thus, the random variable $\sum_{t_1 \le t \le t(R)} \mathbb{1}(g_t \text{ 2-good}) \cdot \mathbb{1}(h_t \text{ 2-good}) \cdot \mathbb{1}(\mathcal{E}_t)$ is stochastically dominated by $A_2 \equiv \text{Bin}\left(n^{0.49}, n^{-0.499}\right)$.

Case 3:

$$\mathbb{Q}(\mathbb{1}(h_t \text{ 3-good}) \cdot \mathbb{1}(\mathcal{E}_t) = 1) \le \frac{33 \cdot n^{7.1/16} + (4e\log_{d-1} n)^2 \cdot n^{1.1\epsilon}}{nd - 2t} \le n^{-0.55}$$

$$\sum_{t_1 \le t \le t(R)} \mathbb{1}(g_t \text{ 3-good}) \cdot \mathbb{1}(\mathcal{E}_t) \le 33 \cdot n^{7.1/16} + (4e\log_{d-1} n)^2 \cdot n^{1.1\epsilon} \le n^{0.45}$$

The random variable

$$\sum_{t_1 \le t \le t(R)} \mathbb{1}(g_t \text{ 3-good}) \cdot \mathbb{1}(\mathcal{E}_t) \cdot \mathbb{1}(h_t \text{ 3-good}) \cdot \mathbb{1}(\mathcal{E}_t) = \sum_{t_1 \le t \le t(R),\ g_t \text{ 3-good},\ \mathcal{E}_t} \mathbb{1}(h_t \text{ 3-good}) \cdot \mathbb{1}(\mathcal{E}_t)$$

is stochastically dominated by $A_3 \equiv \text{Bin}\left(n^{-0.55}, n^{0.45}\right)$.

A_1, A_2, A_3 do not need to be independent. Since $Z_t \le 33 + C_1$ holds for every t in all instances of the joint distribution, We nonetheless have:

$$\mathbb{Q}\left(\sum_{t_1 \le t \le t(R)} Z_t > 630 \cdot (33 + C_1)^2\right) \le \mathbb{Q}(A_1 > 300) + \mathbb{Q}(A_2 > 300) + \mathbb{Q}(A_3 > 30) = o(n^{-2})$$

Now we return to $\sum_{t_1 \le t \le t(R)} Y_t$. Since $\mathcal{E}_t$ is monotone, in any instance, $Z_t \ne Y_t$ for any t implies that $\mathcal{E}_{t(R)}$ does not hold. So $\mathbb{Q}\left(\sum_{t_1 \le t \le t(R)} Z_t \ne \sum_{t_1 \le t \le t(R)} Y_t\right) \le \mathbb{Q}\left(\neg\mathcal{E}_{t(R)}\right) \le \mathbb{Q}\left(U_{t(R)} > 33\right) + \mathbb{Q}\left(L_{t(R)} > (4e\log_{d-1} n)^2\right) + \mathbb{Q}\left(X_{t(R)} > C_1\right)$. So now, we only need to union bound the *end of time* probabilities.

In fact, each of the three probabilities are $o(n^{-2})$, supplied respectively by Lemma 2.13, Lemma 2.12, and Lemma 7.17 below.

Therefore, we have the following upper bound on the probability that at the end of time we need to cut more than $66 + 33 \cdot 630 \cdot (33 + C_1)^2 = C_2$ edges to eliminate $G(\boldsymbol{v})$ from $\mathcal{C}(B_R(\boldsymbol{v}))$:

$$\mathbb{Q}\left(\neg\mathcal{E}_4\right) + \mathbb{Q}\left(\sum_{t_1 \le t \le t(R)} Y_t \ne \sum_{t_1 \le t \le t(R)} Z_t\right) + \mathbb{Q}\left(\sum_{t_1 \le t \le t(R)} Z_t > 630 \cdot (33 + C_1)^2\right) = o(n^{-2})$$

We have proved Proposition 7.1.

Lemma 7.17. $\mathbb{Q}(X_{t(R)+1} > C_1) = o(n^{-2})$.

Proof. To upper bound $X_{t(R)}$, we do a step by step analysis of the BFS that is independent of the Z_t analysis.

Let $w \in V_t(\boldsymbol{v})$. If lower crossing e satisfies $dist_t(w, t_{e_t}) \le 2.2\epsilon\log_{d-1} n$, we call e a *nearby crossing* of w. Define $\hat{X}_t \equiv \max(\mathbb{1}(L_t \le (4e\log_{d-1} n)^2) \cdot X_t, 1)$. In step t, if a new edge is connected, denote it $e_t : g_t \to h_t$.

If $\hat{X}_t < \hat{X}_{t+1}$, then $\hat{X}_{t+1} \ge 2$, and e_t must be one of the three cases below:

(1) e_t is a u-to-v crossing. e_t does not update distance from time t to $t+1$, as we only look at distance in $G_t(\boldsymbol{v}) = G_{t+1}(\boldsymbol{v})$. Let w_{t+1} be the vertex with X_{t+1} nearby crossings at time $t+1$. At time t, without t_{e_t}, w_{t+1} must have $\hat{X}_{t+1} - 1$ nearby crossings. $\hat{X}_{t+1} - 1 \le \hat{X}_t$, so $\hat{X}_{t+1} = \hat{X}_t + 1$.
   - g_t must be attached to a u-vertex.
   - At time t, there exists a crossing e such that $dist_t(h_t, t_e) \le 4.4\epsilon\log_{d-1} n$. This is because $dist_t(w_{t+1}, h_t) \le 2.2\epsilon\log_{d-1} n$, $dist_t(w_{t+1}, t_e) \le 2.2\epsilon\log_{d-1} n$.

- $\hat{X}_{t+1} = \hat{X}_t + 1 \le 2\hat{X}_t$.

(2) e_t is a v-to-u crossing. By a similar analysis,
- At time t, there exists a crossing e such that $dist_t(g_t, t_e) \le 4.4\epsilon \log_{d-1} n$.
- h_t must be attached to a u-vertex.
- $\hat{X}_{t+1} = \hat{X}_t + 1 \le 2\hat{X}_t$.

(3) e_t is not a crossing, so it must update distance. Let the two vertices of e_t be $w_1 \neq w_2$. w_{t+1} (defined in case 1) can be w_1 or w_2. Let $e_1, \cdots, e_k$ be the time-$(t+1)$ nearby lower crossings of w_{t+1}, with $dist_{t+1}(w_{t+1}, t_{e_i}) < dist_t(w_{t+1}, t_{e_i})$, $1 \le i \le k$. Then $\hat{X}_{t+1} - k \le \hat{X}_t$.

For each i, the time-$(t+1)$ shortest path from t_{e_i} to w_{t+1} must pass through e_t. WLOG $dist_t(w_{t+1}, w_1) \le dist_t(w_{t+1}, w_2)$, then $dist_t(w_{t+1}, w_1) = dist_{t+1}(w_{t+1}, w_1)$. The path of each t_{e_i} pass through e_t in the $w_2 \to w_1$ direction, so $dist_t(w_1, t_{e_i}) = dist_{t+1}(w_1, t_{e_i}) < 2.2\epsilon \log_{d-1} n$. Thus, $k \le \hat{X}_t$.
- At least one of g_t, h_t must be attached to w_2, a vertex of time-t distance $\le 2.2\epsilon \log_{d-1} n$ to some t_e. Both g_t, h_t are connected to v-vertices. We get two sub-cases: (3.1) g_t attached to w_2, (3.2) h_t attached to w_2.
- $\hat{X}_{t+1} \le k + \hat{X}_t \le 2\hat{X}_t$.

Note that if $L_t \ge (4e \log_{d-1} n)^2$, $\hat{X}_t = 1 \le \hat{X}_{t-1}$. Thus, in the probability distribution of *all instances of the joint distribution satisfying* $\mathcal{E}_1$, the number of steps t of case 1 is stochastically dominated by $B_1 = \text{Bin}\left(n^{1/2} \log_{d-1} n, \frac{n^{4.4\epsilon}}{nd - 2n^{1/2} \log_{d-1} n}\right)$. The number of steps t of case 2 is stochastically dominated by $B_2 = \text{Bin}\left(n^{4.4\epsilon}, \frac{n^{1/2} \log_{d-1} n}{nd - 2n^{1/2} \log_{d-1} n}\right)$. The numbers of steps t of cases (3.1) and (3.2) are stochastically dominated by $B_3 = \text{Bin}\left(n^{2.2\epsilon}, \frac{n^{1/2} \log_{d-1} n}{nd - 2n^{1/2} \log_{d-1} n}\right)$, $B_4 = \text{Bin}\left(n^{1/2} \log_{d-1} n, \frac{n^{2.2\epsilon}}{nd - 2n^{1/2} \log_{d-1} n}\right)$, respectively. The four binomials do not have to be independent.

Let Z be the number of steps t such that $\hat{X}_t < \hat{X}_{t+1}$. Z is a random variable. Then $\mathbb{Q}(Z > 22) \le \mathbb{P}(B_1 > 6) + \mathbb{P}(B_2 > 6) + \mathbb{P}(B_3 > 5) + \mathbb{P}(B_4 > 5) = o(n^{-2})$. Since $\hat{X}_{t+1} \le 2\hat{X}_t$ in every step in every instance in the joint distribution, $Z \le 22$ at the end of time implies that $\hat{X}_{t(R)+1} \le 2^{22} = C_1$. Therefore,

$$\begin{aligned}\mathbb{Q}(X_{t(R)+1} > C_1) &\le \mathbb{Q}(\hat{X}_{t(R)+1} \neq X_{t(R)+1}) + \mathbb{Q}(\hat{X}_{t(R)+1} > C_1) \\ &\le \mathbb{Q}(L_{t(R)+1} > (4e \log_{d-1} n)^2) + \mathbb{Q}(Z > 22) = o(n^{-2})\end{aligned}$$

□

8. Coupling analysis: eliminate contaminated cycle structure within $G(\boldsymbol{u})$ and $G(\boldsymbol{x})$

In this section, we prove that we can delete 73 edges from $G(\boldsymbol{u}) \cup G(\boldsymbol{x})$ so that the resulting cycle structures of $\boldsymbol{u}$ and $\boldsymbol{x}$ contain no bad edges, w.h.p..

We will analyze the cycle structure of u and x's exploration together. Let $H_t \equiv H_t(\boldsymbol{x}) \cup H_t(\boldsymbol{u})$.

Recall from Lemma 2.7 that the cycle structure of a directed exploration can be built iteratively with single add operations. Call an edge $(u \to w)$ *vertical* if u and w lie on different BFS layers, and *horizontal* otherwise. For $z \in \{u, x\}$ and vertex $w \in V(G_z)$, let its *ancestor path* $\pi_{z \to w}$ be the path from the root z to w consisting of parents (Definition 2.3). After a collision at time t, the current cycle structure is updated with the collision edge (horizontal or vertical) and the endpoints' ancestor paths (vertical).

As $v \in \mathcal{C}$ implies $\text{par}(v) \in \mathcal{C}$, we can equivalently characterize $\mathcal{C}(H_t)$ as the union over collisions of the collision edge plus ancestor paths:

$$\mathcal{C}(H_t) = \bigcup_{\substack{1 \le t \le T \\ h_t \in F_t(z)}} (u_t \to w_t) \cup \pi_{z,u_t} \cup \pi_{z,w_t}$$

— call $(u_t \to w_t) \cup \pi_{z,u_t} \cup \pi_{z,w_t}$ the *contribution* of the collision at time t. If we remove all collisions whose contribution includes at least one bad edge — call this set ERR — the resulting union of contributions, which is the cycle structure of $H_t(\mathbf{z}) \backslash ERR$, contains no bad edges. In the remainder of this section, we show $|ERR| \le 73$ w.h.p..

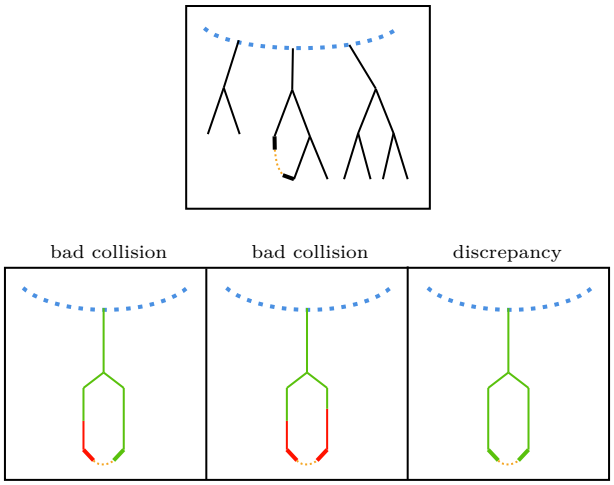

FIGURE 9. A collision contributing bad cycle strucutre is either a discrepancy or bad collision.

Lemma 8.1. *For each collision $g_t \to h_t$ in ERR, either*

(1) *$(g_t \to h_t)$ is a discrepancy, or*
(2) *At least one of $\{g_t, h_t\}$ was bad right before the collision. Call $(g_t \to h_t)$ a bad collision.*

Proof. Edge bad-ness is monotone in ancestor paths: a half-edge whose incident vertex is discovered by a bad half-edge must also be bad. If neither g_t nor h_t are bad right before matching and $(g_t \to h_t) \in ERR$, $(g_t \to h_t)$ must be a discrepancy by construction (see Figures 6, 9). □

For any $1 \le R \le R_{\max}$, let $T = (d-1)^R$. Let I^{disc} be the set of discrepancies, I^{bc} the set of bad collisions, and γ the set of collisions in H_T. By Lemma 8.1,

$$ERR = I^{disc} \cup I^{bc}.$$

We proceed to bound $|I^{disc}|$ and $|I^{bc}|$ individually.

Lemma 8.2 (few discrepancies)**.**

$$\mathbb{Q}\left(|I^{disc}| \ge 5\right) = o(n^{-2})$$

Proof. See proof of Lemma 5.14 of [5]. In [5], $I_{t,2}$ is the event of a discrepancy, and the bound of X in case (1) is precisely our desired result.

□

Next, we bound $|I^{bc}|$. To do so, we first bound the number of bad edges:

Lemma 8.3 (few bad edges)**.** *Let bad_T be the set of all bad half-edges explored up to and including time T.*

$$\mathbb{Q}\left(|\mathsf{bad}_T| \ge 300(\log n)^2 n^{7/16}\right) = o(n^{-2})$$

Proof. Call coupling errors and collisions in G_x *or* G_{uv} source errors. If g is a bad half-edge, it either is involved in a source error or is a descendant of a vertex involved in a source error. Here, b is a descendant of a if $a = \mathrm{par}^k(b)$ for some integer k.

Crucially, all source errors occur after level L_0 given that $\mathcal{E}_0$ holds. Therefore,

$$\begin{aligned}
|\mathsf{bad}_T| &\le 2\left(|I^{\mathrm{disc}}| + |\gamma|\right)(d-1)^{R-L_0} \\
&\le 2\left(|I^{\mathrm{disc}}| + |\gamma|\right)\frac{n^{1/2}\log n}{n^{1/16}} \\
&= 2\left(|I^{\mathrm{disc}}| + |\gamma|\right) n^{7/16}\log n
\end{aligned}$$

where the factor of 2 accounts for the mirroring of errors. By Lemma 2.12,

$$\mathbb{Q}\left(|\gamma(\mathcal{C}(H_t(\boldsymbol{x}))| \geq (2e\log n)^2\right) \leq \exp\{-(\log n)^2\},$$
$$\mathbb{Q}\left(|\gamma(\mathcal{C}(H_t(\boldsymbol{uv}))| \geq (4e\log n)^2\right) \leq \exp\{-(\log n)^2\}.$$

Combining, we can say

$$\mathbb{Q}\left(|\gamma(\mathcal{C}(H_t(\boldsymbol{x}))| + |\gamma(\mathcal{C}(H_t(\boldsymbol{uv}))| \geq 20e^2(\log n)^2\right) \leq \exp\{-(\log n)^2\}$$
$$\implies \mathbb{Q}\left(|\gamma| \geq 149(\log n)^2\right) = o(n^{-2}).$$

By Lemma 8.2,

$$\mathbb{Q}\left(|I^{\text{disc}}| \geq (\log n)^2\right) < \mathbb{Q}\left(|I^{\text{disc}}| \geq 5\right) = o(n^{-2}).$$

Overall,

$$\begin{aligned}\mathbb{Q}\left(|\mathsf{bad}_T| \geq 2\cdot(1+149)(\log n)^2 n^{7/16}\right) &\leq \mathbb{Q}\left(|I^{\text{disc}}| \geq (\log n)^2\right) + \mathbb{Q}\left(|\gamma| \geq 149(\log n)^2\right)\\ &= o\left(n^{-2}\right).\end{aligned}$$

□

Most of the times, there are few bad edges at the end of time, in which the number of bad collisions is dominated by a binomial with constant mean. Only rarely are there too many bad edges. The following lemma formalizes this arugment:

Lemma 8.4 (few bad collisions)**.**

$$\mathbb{Q}\left(|I^{bc}| \geq 68\right) = o(n^{-2})$$

Proof. Define

$$X_{BFS} \equiv \sum_{t=1}^{T} \mathbb{1}((g_t \to h_t) \in I^{bc})$$
$$\mathsf{bad}_{max} \equiv 300(\log n)^2 n^{7/16}$$
$$F_{max} \equiv 2(d-2)(d-1)^{R-1}$$

and note $|I^{bc}| \sim X_{BFS}$. Couple X_{BFS} with random variable X_{SIM}, both over the uniform distribution on $\mathcal{G}$. We define X_{SIM} and the coupling as follows:

Do a BFS exploration from $\delta x \cup \delta u$. At time t $(1 \leq t \leq T)$, sample $U_t \sim \text{Unif}[0,1)$.

- If g_t is bad, let
$$p_l \equiv \frac{F_{t-1}-1}{nd-2t},$$
$$p_r \equiv \frac{F_{max}}{nd-2t}.$$
  - If $U \in [0, p_l)$, increment both X_{BFS} and X_{SIM}. Let h_t be a frontier edge chosen uniformly at random.
  - If $U \in [p_l, p_r)$, increment X_{SIM} only. Sample h_t according to $\mathbb{Q}(h_t | h_t \notin F_{t-1})$.
  - If $U \in [p_r, 1)$, sample h_t according to $\mathbb{Q}(h_t | h_t \notin F_{t-1})$.
- If g_t is good, let
$$p_l \equiv \frac{|\mathsf{bad}_t|}{nd-2t},$$
$$p_r \equiv \frac{\mathsf{bad}_{max}}{nd-2t}.$$
  - If $U \in [0, p_l)$, increment both X_{BFS} and X_{SIM}. Let h_t be a bad edge chosen uniformly at random.
  - If $U \in [p_l, p_r)$, increment X_{SIM} only. Sample h_t according to $\mathbb{Q}(h_t | h_t \notin F_{t-1} \backslash K_{t-1})$.
  - If $U \in [p_r, 1)$, sample h_t according to $\mathbb{Q}(h_t | h_t \notin F_{t-1} \backslash K_{t-1})$.
- If, after timestep t, there are bad_{max} or more bad edges, set $X_{SIM} = \infty$.

The graph we create follows the random regular law, so X_{BFS} is distributed as (2). Furthermore, in any sample, $X_{BFS} \leq X_{SIM}$ as we never increase X_{BFS} without increasing X_{SIM}. Thus, $X_{BFS} \preccurlyeq X_{SIM}$ and

$$\begin{aligned}\mathbb{Q}\left(|I^{bc}| \geq 68\right) &= \mathbb{Q}\left(X_{BFS} \geq 68\right) \\ &\leq \mathbb{Q}\left(X_{SIM} \geq 68\right) \\ &\leq \mathbb{Q}\left(B_1 + B_2 \geq 68\right) \\ &\quad + \mathbb{Q}\left(X_{SIM} = \infty\right)\end{aligned} \tag{3}$$

where we define

$$B_1 \equiv \text{Bin}\left(\mathsf{bad}_{max}, \frac{F_{max}}{nd - 2T}\right),$$

$$B_2 \equiv \text{Bin}\left(T, \frac{\mathsf{bad}_{max}}{nd - 2T}\right).$$

By Chernoff bound,

$$\mathbb{E}[B_2] \leq \mathbb{E}[B_1] \leq \frac{2000(\log n)^3}{dn^{1/16}}.$$

$$\begin{aligned}\mathbb{Q}\left(B_1 \geq 34\right) &\leq \exp\left\{-34 \log\left(\frac{34dn^{1/16}}{2000(\log n)^3}\right)\right\} \\ &\leq \exp\left\{-2 \log\left(\frac{C \cdot n^{17/16}}{(\log n)^{51}}\right)\right\} \\ &= o(n^{-2}).\end{aligned}$$

and the same bound holds for B_2 as it has smaller mean. Thus,

$$\mathbb{Q}(B_1 + B_2 \geq 68) \leq \mathbb{Q}(B_1 \geq 34) + \mathbb{Q}(B_2 \geq 34) = o(n^{-2}).$$

By Lemma 8.3,

$$\mathbb{Q}\left(X_{SIM} = \infty\right) = \mathbb{Q}\left(|\mathsf{bad}_T| \geq \mathsf{bad}_{max}\right) = o(n^{-2}).$$

As both terms in (3) are $o(n^{-2})$, we conclude

$$\mathbb{Q}\left(|I^{bc}| \geq 68\right) = o(n^{-2}).$$

□

Combining Lemmas 8.2 and 8.4,

$$\begin{aligned}\mathbb{Q}\left(|ERR| \geq 73\right) &= \mathbb{Q}\left(|I^{disc}| + |I^{bc}| \geq 73\right) \\ &\leq \mathbb{Q}\left(|I^{disc}| \geq 5\right) + \mathbb{Q}\left(|I^{bc}| \geq 68\right) \\ &= o(n^{-2}).\end{aligned}$$

9. Distinguishablility: concluding the coupling and proving Lemmas 4.4, 4.2

Proposition 9.1. In every instance of the joint distribution, at the end of time, if $dist(\mathcal{C}(B_R(\boldsymbol{x})), \mathcal{C}(B_R(\boldsymbol{u}))) \leq C_2 + 73$, there exist no more than $C_2 + 73$ delete (or add) operations, each removing (or adding) at most $2R$ edges in the cycle structure, that turn $\mathcal{C}(B_R(\boldsymbol{x}))$ into $\mathcal{C}(B_R(\boldsymbol{u}))$.

Proof. In the previous sections we bound $dist(\mathcal{C}(B_R(\boldsymbol{x})), \mathcal{C}(B_R(\boldsymbol{u})))$ by deleting edges in $\mathcal{C}(B_R(\boldsymbol{u}))$ and $\mathcal{C}(B_R(\boldsymbol{x}))$ until the remaining cycle structure is that of $\mathcal{C}(K_{t(R)}(\boldsymbol{x}))$. Because operations are reversible, we only need to look at delete operations in both graphs. We will prove that, in both graphs, we can delete in such a way that removes at most $2R$ edges in each deletion. In fact, if in each step we delete the edge e that is not in $\mathcal{C}(K_{t(R)}(\boldsymbol{x}))$, and that one of e's vertices has the maximum depth, and then recursively prune the leaves, the maximum depth of what remains of the cycle structure does not exceed R (otherwise, the new deepest vertex has its depth changed by the operation, so before the operation it would have greater depth than e's incident vertices). Therefore, in each delete operation we remove no more than $2R$ edges. □

Combining two previous sections, we obtained that $\mathbb{Q}(dist(\mathcal{C}(B_R(\boldsymbol{x})),\mathcal{C}(B_R(\boldsymbol{u}))) > C_2 + 73) = o(n^{-2})$. Similarly, $\mathbb{Q}(dist(\mathcal{C}(B_R(\boldsymbol{y})),\mathcal{C}(B_R(\boldsymbol{v}))) > C_2 + 73) = o(n^{-2})$. Let $C_3 \equiv 2C_2 + 2\cdot 73$, and $\mathcal{E}_5$ denote the event that both distances $\le C_2 + 73$ in the joint distribution. Then $\mathbb{Q}(\text{not } \mathcal{E}_5) = o(n^{-2})$. The goal is to prove Lemma 4.4, that is, $\mathbb{P}(\mathcal{E} \text{ and not } \mathcal{E}_0(\boldsymbol{u},\boldsymbol{v})) = o(n^{-2})$.

$$\mathbb{P}(\mathcal{E} \text{ and not } \mathcal{E}_0(\boldsymbol{u},\boldsymbol{v})) = \mathbb{Q}(\mathcal{E} \text{ and } B_R(\boldsymbol{u}) \cong B_R(\boldsymbol{v})) \le \mathbb{Q}(\text{not } \mathcal{E}_5) + \mathbb{Q}(\mathcal{E}, B_R(\boldsymbol{u}) \cong B_R(\boldsymbol{v}), \text{ and } \mathcal{E}_5)$$

It remains to prove that the second $\mathbb{Q}$ is $o(n^{-2})$. By Lemma 4.5, $\mathcal{E}$ and $dist(\mathcal{C}(B_R(\boldsymbol{x})),\mathcal{C}(B_R(\boldsymbol{u}))) \le C_2+73$ implies that the neighborhood structure of $\boldsymbol{x}$, $\mathcal{T}_R(\boldsymbol{x})$, is in $\boldsymbol{\Omega}_{dir,R}$.

Further, let $\boldsymbol{\Xi}$ denote the subset of possible cycle structures $\mathcal{C}$ for which $\gamma(\mathcal{C}) \le (2e\log n)^2$, and note that $\mathbb{Q}\left(\mathcal{C}(B_R(\boldsymbol{y})) \notin \boldsymbol{\Xi}\right) = o(n^{-2})$ by Lemma 2.12.

By Proposition 9.1, $B_R(\boldsymbol{u}) \cong B_R(\boldsymbol{v})$ and $\mathcal{E}_5$ implies that we can turn $\mathcal{C}(B_R(\boldsymbol{y}))$ into $\mathcal{C}(B_R(\boldsymbol{x}))$ with no more than C_3 add/delete operations, where each operation adds/removes at most $2R$ edges. Let this be event $\mathcal{E}_6$. Altogether this gives

$$\mathbb{Q}(\mathcal{E}, B_R(\boldsymbol{u}) \cong B_R(\boldsymbol{v}), \text{ and } \mathcal{E}_5) \le o(n^{-2}) + \mathbb{Q}(\mathcal{T}_R(\boldsymbol{x}) \in \boldsymbol{\Omega}_{dir,R}, \mathcal{C}_R(\boldsymbol{y}) \in \boldsymbol{\Xi}, \mathcal{E}_6)$$

Since the $\boldsymbol{x}$- and $\boldsymbol{y}$-explorations are independent under $\mathbb{Q}$, the RHS is upper bounded by

$$o_n(n^{-2}) + \sum_{\mathcal{C}\in\boldsymbol{\Xi}} \mathbb{Q}\left(\mathcal{C}(B_R(\boldsymbol{y})) = \mathcal{C}\right) \mathbb{Q}(\mathcal{T}_R(\boldsymbol{x}) \in \boldsymbol{\Omega}_{dir,R}, \mathcal{E}_6) \tag{9.1}$$

It follows from Lemma 4.6, taking $\rho = 3$, that if $R \ge R_+(\Delta)$ with Δ sufficiently large, then

$$\max\left\{\mathbb{Q}\left(\mathcal{T}_R(\boldsymbol{x}) = \mathcal{T}'\right) : \mathcal{T}' \in \boldsymbol{\Omega}_{dir,R}\right\} \ll n^{-3}.$$

There is a bijection between neighborhood types and cycle structures. For $\mathcal{C} = \mathcal{C}(B_R(\boldsymbol{y})) \in \boldsymbol{\Xi}$, the number of $\mathcal{C}(B_R(\boldsymbol{x}))$ that satisfies $\mathcal{E}_6$ is, crudely, at most $(\log n)^{8C_3}$: first, the number of edges in $\mathcal{C}$ is $\le 2R\gamma(\mathcal{C}) \le (\log n)^{3.1}$. Each operation can increase the total number of edges by at most $2R \le 2R_{\max} \le (\log n)^{1.1}$, so during C_3 add/delete operations the total number of edges cannot increase beyond $(\log n)^{3.2}$. Recalling Definition 2.10, for each add operation it suffices to either specify the start point, end point, length of the new segment, or, the start point, length of segment, length of hanging cycle. For each delete operation it suffices to specify a single cut edge. Note that due to $\mathcal{E}_6$, the length of the new path and hanging cycle cannot exceed $2R_{\max}$. So number of possible operations (including the identity operation) at each step is then $\le 3(\log n)^{3.2+3.2+1.1} + 1 \le (\log n)^8$. Given any $\mathcal{C} \in \boldsymbol{\Xi}$, the number of distinct $\mathcal{C}(B_R(\boldsymbol{x}))$ satisfying $\mathcal{E}_6$ is bounded by $\le (\log n)^{8C_3}$ as claimed. Then altogether

$$\mathbb{Q}(\mathcal{T}_R(\boldsymbol{x}) \in \boldsymbol{\Omega}_{dir,R}, \mathcal{E}_6) \le \frac{(\log n)^{8C_3}}{n^3} \ll n^{-2}$$

Substituting into (9.1) proves Lemma 4.4.

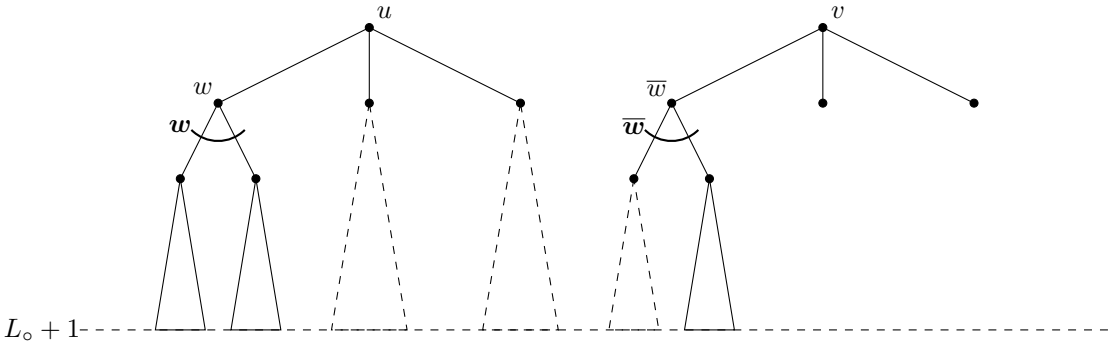

FIGURE 10. Case 1: Solid triangles represent trees that do not intersect the rest of the graph, while dotted triangles represent descendant subgraphs that may not be trees, and may intersect the rest of the graph. If there is a direction $\overline{\boldsymbol{w}}_1 \subseteq \overline{\boldsymbol{w}}$ of size $|\overline{\boldsymbol{w}}_1| = d-2$ such that $B_{L_\circ}(\overline{\boldsymbol{w}}_1) \cap B_{L_\circ}(\boldsymbol{w}) = \emptyset$, then $B_R(\overline{\boldsymbol{w}}_1)$ cannot be isomorphic to $B_R(\varphi^{-1}(\overline{\boldsymbol{w}}_1))$. Thus $B_R(u) \ncong B_R(v)$. Note that $\overline{\boldsymbol{w}}$ does not have to be the $d-1$ downward directions, as shown in the figure. The cases where $\overline{\boldsymbol{w}}$ contains v bear the same analysis.

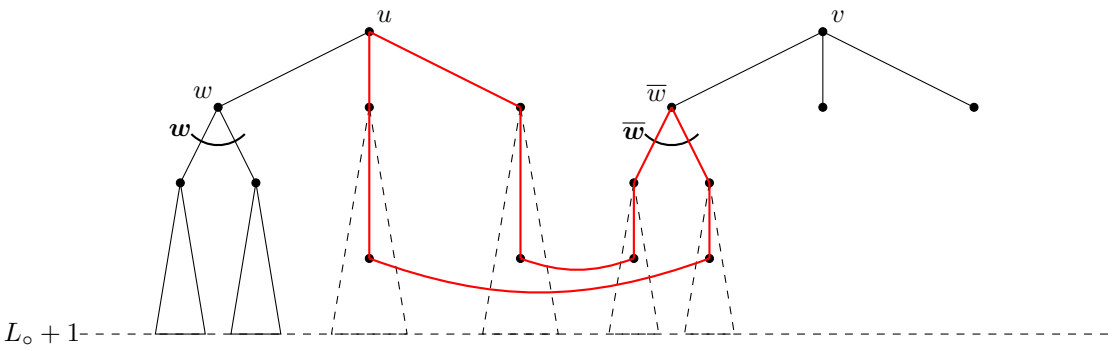

FIGURE 11. Case 2: If there is not a good direction $\overline{\boldsymbol{w}}_1 \subseteq \overline{\boldsymbol{w}}$ of size $|\overline{\boldsymbol{w}}_1| = d-2$, $B_{L_\circ}(\overline{\boldsymbol{w}})$ is not a tree. Thus $B_R(\boldsymbol{w}) \ncong B_R(\overline{\boldsymbol{w}})$, and so $B_R(u) \ncong B_R(v)$.

9.1. **Proof of Lemma 4.2.**

Proof. Given a G that satisfies (1) (2), let $u \neq v \in V$. We only consider the subgraph $H = B_{L_\circ+1}(\delta u) \cap B_{L_\circ+1}(\delta v)$. By (1) the total number of within $B_{L_\circ+1}(\delta u)$ collision(s), within $B_{L_\circ+1}(\delta v)$ collision(s), and $u-v$ crossing(s) combined does not exceed 2. For each of such collision/crossing edge, composed of two half-edges g, h, trace the ancestors of g, h until we reach the half-edges g', h' attached to u or v (g' can be g itself). This gives us at most 4 half-edges attached to either u or v. And there are at least $2 \cdot 3 - 4 = 2$ remaining half-edges attached to u or v, whose descendants do not contribute to a collision or crossing. Let one of them be e, which matches half-edge f attached to vertex w. $w \notin \{u, v\}$ since e does not lead to a self-loop or crossing.

WLOG e is attached to u. Let $\boldsymbol{w} = \delta w \backslash f$ be the $d-1$ downward directions of w. No vertex in $B_{L_\circ}(\boldsymbol{w})$ is part of a collision or crossing. No vertex in $B_{L_\circ}(\boldsymbol{w})$ has an edge in H with a v-vertex (so $B_{L_\circ}(\boldsymbol{w}) \cap V(v) = \varnothing$), or with a u-vertex that descends from an half-edge in δu that is not e. Any path contained in H, that starts from a v-vertex and reaches a vertex in $B_{L_\circ}(\boldsymbol{w})$, must enter $B_{L_\circ}(\boldsymbol{w})$ through edge $u-w$. $B_{L_\circ}(\boldsymbol{w})$ is a tree, isolated from the rest of $B_{L_\circ+1}(\delta u) \cup B_{L_\circ+1}(\delta v)$ except for edge $u-w$.

We prove by contradiction that $B_{R+1}(\delta u) \ncong B_{R+1}(\delta v)$. Let there be an isomorphism $\phi(B_{R+1}(\delta u)) = B_{R+1}(\delta v)$. Let $\overline{w} = \phi(w)$, then $\overline{w}$ is distance 1 from v, so $\overline{w} \neq w$. Let $\overline{f} = \phi(f)$, $\overline{\boldsymbol{w}} = \delta \overline{w} \backslash \overline{f} = \phi(\boldsymbol{w})$. Then $B_R(\boldsymbol{w}) \cong B_R(\overline{\boldsymbol{w}})$; $B_{L_\circ}(\overline{\boldsymbol{w}})$ is a tree. $B_{L_\circ}(\overline{\boldsymbol{w}})$ must be one of the following two cases (see Figures 10, 11):

Case 1: if there is a direction $\overline{\boldsymbol{w}}_{\mathbf{1}} \subset \overline{\boldsymbol{w}}$, $|\overline{\boldsymbol{w}}_{\mathbf{1}}| = d-2$, such that $B_{L_\circ}(\boldsymbol{w}) \cap B_{L_\circ}(\overline{\boldsymbol{w}}_{\mathbf{1}}) = \varnothing$, then directions $\overline{\boldsymbol{w}}_{\mathbf{1}}$ and $\phi^{-1}(\overline{\boldsymbol{w}}_{\mathbf{1}})$ satisfy (a). This is because $\phi^{-1}(\overline{\boldsymbol{w}}_{\mathbf{1}}) \subset \phi^{-1}(\overline{\boldsymbol{w}}) = \boldsymbol{w}$, and $B_{L_\circ}(\phi^{-1}(\overline{\boldsymbol{w}}_{\mathbf{1}})), B_{L_\circ}(\overline{\boldsymbol{w}}_{\mathbf{1}})$ are both trees. By (a), (b) holds, that is, $B_R(\phi^{-1}(\overline{\boldsymbol{w}}_{\mathbf{1}})) \ncong B_R(\overline{\boldsymbol{w}}_{\mathbf{1}})$. Contradiction.

Case 2: for any half-edge $k \in \overline{\boldsymbol{w}}$, $B_{L_\circ}(\boldsymbol{w}) \cup B_{L_\circ}(\overline{\boldsymbol{w}} \backslash k) \neq \varnothing$. So there is a non-self-intersecting path in $B_{L_\circ}(\overline{\boldsymbol{w}})$ that does not contain k, starting at $\overline{w}$ and ending at a vertex in $B_{L_\circ}(\boldsymbol{w})$: the latter implies that the path must go through edge $u-w$. Pick an arbitrary $k \in \overline{\boldsymbol{w}}$ and let the path be $\pi_1 \subset B_{L_\circ}(\overline{\boldsymbol{w}})$. π_1 starts with $k_1 \subset \overline{\boldsymbol{w}}$. Given $B_{L_\circ}(\boldsymbol{w}) \cap B_{L_\circ}(\overline{\boldsymbol{w}} \backslash k_1) \neq \varnothing$, there is another non-self-intersecting path π_2 that does not contain k_1 and starts with $k_2 \subset \overline{\boldsymbol{w}}$. Thus, π_1, π_2 depart $\overline{\boldsymbol{w}}$ in two different edges. Both paths go through edge $u-w$. Trace π_1, π_2 from $\overline{w}$ to w. Let x be the first vertex, other than the start $\overline{w}$, shared by both paths. x exists because u is shared by both paths. Moreover, $x \neq \overline{w}$ because neither path self-intersects. The sections of π_1, π_2 respectively from $\overline{w}$ to x do not share any edge, as $k_1 \neq k_2$. Walking from $\overline{w}$ to x along π_1, and back to $\overline{w}$ along π_2^{-1}, forms a cycle in $B_{L_\circ}(\overline{\boldsymbol{w}})$. This contradicts the isomorphism which requires $B_{L_\circ}(\overline{\boldsymbol{w}})$ to be a tree.

As a result, $B_{R+1}(u) \ncong B_{R+1}(v)$ for any $u \neq v \in V$. We can reconstruct G. □

REFERENCES

1. Béla Bollobás, *Distinguishing vertices of random graphs*, Graph Theory (Béla Bollobás, ed.), North-Holland Mathematics Studies, vol. 62, North-Holland, 1982, pp. 33–49.
2. Frank Harary, *A survey of the reconstruction conjecture*, Graphs and Combinatorics (Berlin, Heidelberg) (Ruth A. Bari and Frank Harary, eds.), Springer Berlin Heidelberg, 1974, pp. 18–28.
3. Paul J. Kelly, *A congruence theorem for trees.*, Pacific Journal of Mathematics **7** (1957), no. 1, 961 – 968.

4. Elchanan Mossel and Nathan Ross, *Shotgun assembly of labeled graphs*, IEEE Transactions on Network Science and Engineering **6** (2019), no. 2, 145–157.
5. Elchanan Mossel and Nike Sun, *Shotgun assembly of random regular graphs*, arXiv preprint arXiv:1512.08473 (2015).
6. Nicholas C Wormald et al., *Models of random regular graphs*, London mathematical society lecture note series (1999), 239–298.

Appendix A

Here we reproduce Lemma 7.5 and prove it.

Lemma A.1 (Lemma 7.5)**.** *Let N be a graph rooted in t where each vertex is of distance $\leq r$ from the root. Then if edge e is in some cycle C, there must be some cycle C' with length $\leq 2r+1$ such that $e \in C'$.*

Proof. Given the fixed H, we do any BFS exploration starting from v. Let C be the shortest cycle containing e, and the vertices of e are s, t with $depth(s) = dist(s, v) \leq dist(t, v) = depth(t)$.

If e is a collision edge in the BFS, then consider the BFS upward path from s to v, and the BFS upward path from t to v. The two paths first intersect at w (w could be v). Then w, s, t, w is a cycle of length $\leq dist(s, v) + 1 + dist(t, v) \leq 2d + 1$.

If e isn't a collision, then $dist(s, v) < dist(t, v)$. We travel along C in the direction of $s \to t$. Let $x \to y$ be the first edge we find on C such that $dist(x, v) \geq dist(y, v)$ (the first non-downward edge; x, y exist as $d < \infty$; x can be t). Let π be the BFS upward path from y to v. Then $x - y \notin \pi$. Also, if $e \in \pi$, it must be $t \to s$ in π. Then we can replace the $t \to x \to y$ part of C with the shorter $y \to t$ part of π, forming a shorter cycle containing e: absurd. Thus, π does not contain e. Consider the path $x \to t \to s \to v$ and the path $(x \to y) \cup \pi$. Remove edges above their first intersection (e won't be removed). This is a $\leq 2d+1$ cycle C' with $e \in C'$. □