THE DIAMETER OF INHOMOGENEOUS RANDOM GRAPHS

NICOLAS FRAIMAN AND DIETER MITSCHE

Abstract. In this paper we study the diameter of inhomogeneous random graphs $G(n, \kappa, p)$ that are induced by irreducible kernels κ. The kernels we consider act on separable metric spaces and are almost everywhere continuous. We generalize results known for the Erdős-Rényi model $G(n, p)$ for several ranges of p. We find upper and lower bounds for the diameter of $G(n, \kappa, p)$ in terms of the expansion factor and two explicit constants that depend on the behavior of the kernel over partitions of the metric space.

1. Introduction

In this work we study metric properties of inhomogeneous random graphs, where edges are present independently but with unequal edge occupation probabilities. We study the behavior of the diameter for different ranges of the mean edge density. Under weak assumptions we find tight asymptotic bounds of the diameter for connected graphs in this random graph model.

Let $\mathcal{S}$ be a separable metric space and μ a Borel probability measure on $\mathcal{S}$. Let $\kappa : \mathcal{S} \times \mathcal{S} \to [0, 1]$ be a measurable symmetric kernel. The *inhomogeneous random graph* with kernel κ and density parameter p (depending on n) is the random graph $G(n, \kappa, p) = (V, E)$ where the vertex set is $V = \{1, \ldots, n\}$ and we connect each pair of vertices $i, j \in V$ independently with probability $p_{ij} = \kappa(X_i, X_j)p$, where $X_1, \ldots, X_n$ are independent μ-distributed random variables on $\mathcal{S}$.

We study the asymptotic expansions for distances in the graph $G(n, \kappa, p)$ by associating to $G(n, \kappa, p)$ two graphs induced by the kernel κ. Given two subsets $\mathcal{A}, \mathcal{B} \subset \mathcal{S}$, let

$$\kappa_\ell(\mathcal{A}, \mathcal{B}) = \operatorname{ess\,inf}\{\kappa(x, y) : x \in \mathcal{A}, y \in \mathcal{B}\},$$
$$\kappa_u(\mathcal{A}, \mathcal{B}) = \operatorname{ess\,sup}\{\kappa(x, y) : x \in \mathcal{A}, y \in \mathcal{B}\}.$$

A finite partition $\mathbb{A} = \{\mathcal{A}_1, \ldots, \mathcal{A}_m\}$ of $\mathcal{S}$ is called *essential*, if it has no measure zero sets and it covers all of $\mathcal{S}$ except possibly a measure zero set, that is $\mu(\mathcal{A}_i) > 0$ and $\mu(\mathcal{S} \setminus \cup_{i=1}^m \mathcal{A}_i) = 0$. All partitions considered in the manuscript are essential partitions. For a finite partition $\mathbb{A} = \{\mathcal{A}_1, \ldots, \mathcal{A}_m\}$ of $\mathcal{S}$, we define the *lower partition graph* $P_\ell(\mathbb{A})$ induced by $\mathbb{A}$ as the graph with vertex set $\mathbb{A}$ and where $(\mathcal{A}_i, \mathcal{A}_j)$ is an edge if $\kappa_\ell(\mathcal{A}_i, \mathcal{A}_j) > 0$. Analogously, we define the *upper partition graph* $P_u(\mathbb{A})$ as

Date: January 23, 2018.
2010 *Mathematics Subject Classification.* 60C05, 05C80.
Key words and phrases. random graphs, diameter, concentration inequalities, neighborhood expansion.
N.F. was partially supported by an NSERC Postdoctoral Fellowship and the Department of Mathematics at Harvard University.

the graph with vertex set $\mathbb{A}$ and where $(\mathcal{A}_i, \mathcal{A}_j)$ is an edge if $\kappa_u(\mathcal{A}_i, \mathcal{A}_j) > 0$. Note that in both graphs we allow loops if $\kappa_\ell(\mathcal{A}_i, \mathcal{A}_i) > 0$ ($\kappa_u(\mathcal{A}_i, \mathcal{A}_i) > 0$, respectively). We only consider finite partitions with at least three parts (we see below that we need this assumption for a claim about the diameter of the lower partition graph when refining a partition).

A kernel κ on $(\mathcal{S}, \mu)$ is *reducible* if there exists a set $\mathcal{A} \subset \mathcal{S}$ with $0 < \mu(\mathcal{A}) < 1$ such that $\kappa = 0$ almost everywhere on $\mathcal{A} \times (\mathcal{S} \setminus \mathcal{A})$. Otherwise κ is *irreducible*.

Throughout the paper we always assume that n is sufficiently large. We say that a sequence of events holds *with high probability*, if it holds with probability tending to 1 as $n \to \infty$. Since we are interested in results that hold with high probability (rather than almost surely), we can work with essential partitions. Throughout the paper we denote by ω an arbitrary function tending to infinity with n; however, in all statements below the statement is stronger (the conditions are less restrictive, respectively), if $\omega \to \infty$ at a slower speed.

For two vertices $u, v \in V$ belonging to the same connected component of a graph $G = (V, E)$, denote by $d_G(u, v)$ the graph distance between u and v, that is, the number of edges on a shortest path between them. For a connected graph G, let $\operatorname{diam} G = \max_{u,v} d_G(u, v)$. We study the diameter of $G(n, \kappa, p)$ by studying the diameters of the induced graphs $P_\ell(\mathbb{A})$ and $P_u(\mathbb{A})$.

We define the following two constants:

$$\Delta_\ell := \inf_{\mathbb{A}} \operatorname{diam} P_\ell(\mathbb{A}) \qquad \text{and} \qquad \Delta_u := \sup_{\mathbb{A}} \operatorname{diam} P_u(\mathbb{A}),$$

where $\mathbb{A}$ ranges over all partitions with no measure zero sets. Next, we define the *expansion factor*

$$\Phi := \left\lceil \frac{\log n}{\log np} \right\rceil .$$

This quantity is about the diameter of $G(n, p)$, where $G(n, p)$ is the Erdős-Rényi graph, as first shown by [5]. In order to simplify the statements of our results, we consider values of n and p for which

$$(1) \qquad \frac{(np)^\Phi}{n} - \omega \log n \to \infty \qquad \text{and} \qquad \frac{(np)^{\Phi-1}}{n} - \frac{1}{\omega} \log n \to -\infty,$$

In fact, there is no need to have the same ω in both conditions, but since both conditions of (1) are less restrictive if $\omega \to \infty$ at a slower rate, we may as well assume that it is the same ω. For the relation of Φ to the diameter of the Erdős and Rényi model $G(n, p)$, recall the following lemma, first proved by [5]:

Lemma 1 ([6], Corollary 10.12)**.** *Let $np \geq \omega(\log n)^3$ and let $k = k(n) \in \mathbb{N}$. Assume $(np)^k/n - 2\log n \to \infty$ and $(np)^{k-1}/n - 2\log n \to -\infty$. Then, with high probability* $\operatorname{diam}(G(n, p)) = k$.

Observe that if $np \geq \omega(\log n)^3$, and n and p satisfy the conditions given in (1), then the assumptions of Lemma 1 are satisfied. In this case, for k as in Lemma 1, clearly $k = \Phi$. (Our results hold also for the ranges in between with the obvious changes. Since this is handled as in $G(n, p)$, we focus on the one–value case for the sake of clarity).

Throughout the paper we assume the following regularity conditions hold:

Regularity conditions (R). Let $G \in G(n, \kappa, p)$ with κ being irreducible, continuous $(\mu \otimes \mu)$-almost everywhere. Moreover, $np \geq \omega(\log n)^3$, and we assume that the hypotheses given in (1) hold. We also assume that $p < 1$, so that with high probability $\Phi \geq 2$. We also consider only kernels so that $\Delta_u \geq 2$.

In Section 3 we show our first main result:

Theorem 1. *Suppose that regularity conditions (R) hold and suppose $\Delta_\ell < \infty$. Then*

$$\Delta_u \leq \Delta_\ell \leq \Delta_u + 2,$$

both bounds being tight.

Remark 1. *The assumption $\Delta_\ell < \infty$ is needed to get the desired inequalities: For instance, consider the case $\mathcal{S} = [0,1]$ with Lebesgue measure μ and for $x \leq y$, let $\kappa(x,y) = f(x)$ for some function f that tends to 0 as $x \to 0$ (for $x > y$, κ here and below is extended by symmetry). Then for any finite partition $\mathbb{A}$, we have $\operatorname{diam} P_u(\mathbb{A}) = 1$ and $\operatorname{diam} P_\ell(\mathbb{A}) = \infty$, since the part containing 0 is not connected to any other part in P_ℓ, but the part containing 0 as well as all other parts are connected between themselves by an edge in P_ℓ, and hence $\Delta_\ell = \infty$ whereas $\Delta_u = 1$.*

Our second main result is the following:

Theorem 2. *Suppose that regularity conditions (R) hold. With high probability, the following statements hold:*

(i) If $\Phi < \Delta_u$, then

$$\Delta_u \leq \operatorname{diam} G(n, \kappa, p) \leq \Delta_\ell.$$

(ii) If $\Delta_u \leq \Phi < \Delta_\ell$, then

$$\Phi \leq \operatorname{diam} G(n, \kappa, p) \leq \Delta_\ell.$$

Moreover, if there exists a partition $\mathbb{A}$ and $\mathcal{A}_i$ and $\mathcal{A}_j$, with no walk of length exactly Φ between them in $P_u(\mathbb{A})$, then

$$\Phi + 1 \leq \operatorname{diam} G(n, \kappa, p).$$

(iii) If $3 \leq \Delta_\ell \leq \Phi = O(1)$, then

$$\Phi \leq \operatorname{diam} G(n, \kappa, p) \leq \Phi + 1.$$

Moreover,

$$\operatorname{diam} G(n, \kappa, p) = \Phi + 1,$$

if and only if there exists a partition $\mathbb{A}$ and $\mathcal{A}_i$ and $\mathcal{A}_j$, with no walk of length exactly Φ between them in $P_u(\mathbb{A})$.

(iv) If $3 \leq \Delta_\ell < \infty$ and $\Phi = \omega(1)$, then

$$\Phi \leq \operatorname{diam} G(n, \kappa, p) \leq \Phi(1 + o(1)).$$

Remark 2. *Continuing Remark 1, if $\Delta_u = 1$ and $\Delta_\ell = \infty$, the value of p can be chosen so that any value in $[2, \infty)$ is attained for $\operatorname{diam} G(n, \kappa, p)$: indeed, using the notation of Remark 1, if for $x \leq y$, we have say $f(x) = x^{1/x}$ (note that κ is bounded, irreducible and continuous, $f(x)$ is monotonically increasing on $[0,1]$, and $f(x) \to 0$ as $x \to 0$), then with high probability we have $\operatorname{diam} G(n, \kappa, p) = \infty$: indeed, first note that with high probability we find one vertex u such that $X_u \leq C \log n / n$ for C being large enough. Next, such a vertex is isolated, since the expected number* *of neighbors of such a vertex is at most* $n(C\log n/n)^{n/C\log n} \to 0$*, and the result follows from Markov's inequality. On the other hand, if for* $x \le y$*,* $f(x) = x^{1/2-\varepsilon}$ *for some small* $\varepsilon > 0$ *(again* κ *is irreducible, continuous and bounded, and again* $f(x)$ *is monotonically increasing on* $[0,1]$*, and* $f(x) \to 0$ *as* $x \to 0$*) and* $p = \Theta(1)$*, then with high probability* $\operatorname{diam} G(n,\kappa,p) = 2$*: to see this, first note that with high probability for all* u*,* $X_u \ge 1/n\log n$*. Next, with high probability all vertices are at distance* 2 *from each other; indeed, if for all* u*,* $X_u \ge 1/n\log n$*, for a fixed pair of vertices, the probability of having no common neighbor is at most*

$$\Big(1 - (1/(n\log n))^{2(\frac{1}{2}-\varepsilon)}p\Big)^{n-2} \le e^{-n^{2\varepsilon+o(1)}},$$

and by a union bound over all $\binom{n}{2}$ *pairs of vertices, we see that with high probability each pair of vertices has a common neighbor. By suitable choices of* p *any value in* $[2,\infty)$ *can be obtained (the same example can be used to show that all pairs of vertices are at distance at most* k *for any* $k \ge 3$*, and some are actually at distance exactly* k*).*

1.1. **Background and history.** A discrete version of this model was introduced by Söderberg [19]. The sparse case (when the number of edges is linear in the number n of vertices) was studied in detail by Bollobás, Janson and Riordan [7]. Among other things they give an asymptotic formula for the diameter of the giant component when it exists. Connectivity at the intermediate case was analyzed by Devroye and Fraiman [11]. The dense case (when the number of edges is quadratic in n) is closely related with the theory of graph limits started by Lovász and Szegedy [15] and further studied in depth by Borgs, Chayes, Lovász, Sós and Vesztergombi [8, 9] among others. For a thorough introduction to the subject of graph limits see the book by Lovász [14]. Recently, the authors of [12] studied the clique number of dense inhomogeneous graphs.

The diameter of random graphs has been studied widely. In particular, for the Erdős-Rényi model, Bollobás [5] generalized the results of Klee and Larman [13] characterizing the case of constant diameter. Later, Chung and Lu [10] proved concentration results in various different ranges. More recently, Riordan and Wormald [18] completed the program to study the missing cases for the Erdős-Rényi model.

The critical window, when $p = 1/n + cn^{-4/3}$, for $G(n,p)$ is much harder to analyze. Nachmias and Peres [16] obtained the order of the diameter, namely $n^{1/3}$. Addario–Berry, Broutin and Goldschmidt [1, 2] proved convergence, in the Gromov–Hausdorff distance, of the rescaled connected components to a sequence of continuous compact metric spaces. In particular, the diameter rescaled by $n^{-1/3}$ converges in distribution to an absolutely continuous random variable with finite mean. Their approach was extended by Bhamidi, Sen and Wang [4] to the Norros–Reittu [17] random graph model, and then further generalized by Bhamidi, Broutin, Sen and Wang [3].

1.2. **Structure of the paper.** In Section 2 we introduce all concepts, additional definitions and results needed to prove Theorem 2. In Section 3 we prove Theorem 1 on the behavior of the upper and lower diameters Δ_u and Δ_ℓ. In Section 4 we prove that the number of vertices that are at a fixed distance from a given vertex grows exponentially as a function of the distance. Finally, in Section 5 we combine the results of the previous sections to give the proof of Theorem 2.

2. Framework

In this paper we follow the notation from [7] with minor changes. We also use the following standard notation: for functions $f(n)$ and $g(n)$ we write $f = O(g)$ if $|f|/|g|$ is bounded and $f = o(g)$ if $|f|/|g| \to 0$. We say that $f = \Omega(g)$ if $g = O(f)$, and $f = \Theta(g)$ if both $f = O(g)$ and $f = \Omega(g)$ holds.

Given a subset $\mathcal{A} \subset \mathcal{S}$ we write $V(\mathcal{A})$ for the set of vertices with type in $\mathcal{A}$, i.e.,

$$V(\mathcal{A}) = \left\{ v \in V : X_v \in \mathcal{A} \right\}.$$

The asymptotic expansions for distances in the graph $G(n, \kappa, p)$ are obtained by looking at the lower and upper partition graphs $P_\ell(\mathbb{A})$ and $P_u(\mathbb{A})$ of a finite partition $\mathbb{A} = \{\mathcal{A}_1, \ldots, \mathcal{A}_m\}$ of $\mathcal{S}$, as defined in the introduction. These graphs are finite graphs that describe approximations of κ that may be successively refined. More formally, we have the following definition:

Definition 1. We say that a partition $\mathbb{A}$ is a *refinement* of $\mathbb{B}$, denoted by $\mathbb{A} \prec \mathbb{B}$, if for every $\mathcal{A} \in \mathbb{A}$ there exists $\mathcal{B} \in \mathbb{B}$ such that $\mathcal{A} \subset \mathcal{B}$. Note that in this case, for each $\mathcal{B}_i \in \mathbb{B}$ there exists $m_i \in \mathbb{N}$ and $\mathcal{A}_1^{(i)}, \ldots, \mathcal{A}_{m_i}^{(i)} \in \mathbb{A}$ such that $\mathcal{B}_i = \cup_{s=1}^{m_i} \mathcal{A}_s^{(i)}$ μ-almost everywhere.

Let us examine the effect of a refinement on the partition graphs: it is clear that $\kappa_u(\mathcal{B}_i, \mathcal{B}_j) > 0$ if and only if there exist $\mathcal{A}_i, \mathcal{A}_j$ with $\mathcal{A}_i \subset \mathcal{B}_i$ and $\mathcal{A}_j \subset \mathcal{B}_j$ such that $\kappa_u(\mathcal{A}_i, \mathcal{A}_j) > 0$. This implies that $P_u(\mathbb{B})$ is obtained from $P_u(\mathbb{A})$ by contracting the vertices $\mathcal{A}_i \subset \mathcal{B}_i$ into one vertex $\mathcal{B}_i$. In particular,

$$\operatorname{diam} P_u(\mathbb{B}) \le \operatorname{diam} P_u(\mathbb{A}).$$

On the other hand, $\kappa_\ell(\mathcal{B}_i, \mathcal{B}_j) > 0$ if and only if for all $\mathcal{A}_i \subset \mathcal{B}_i$ and $\mathcal{A}_j \subset \mathcal{B}_j$ we have $\kappa_\ell(\mathcal{A}_i, \mathcal{A}_j) > 0$. This implies that the graph obtained by splitting each vertex $\mathcal{B}_i$ into the parts of $\mathbb{A}$ that it contains is a subgraph of $P_\ell(\mathbb{A})$. In particular,

$$\operatorname{diam} P_\ell(\mathbb{B}) \ge \operatorname{diam} P_\ell(\mathbb{A}).$$

(If we had not made our assumption on considering partitions with at least three parts, then the claim could be false when splitting a partition with one or two parts). Note also that if $\mathbb{A} \prec \mathbb{B}$ and $\mathcal{A}_i \subset \mathcal{B}_i$ and $\mathcal{A}_j \subset \mathcal{B}_j$, we always have

$$\begin{aligned} d_{P_\ell(\mathbb{A})}(\mathcal{A}_i, \mathcal{A}_j) &\le d_{P_\ell(\mathbb{B})}(\mathcal{B}_i, \mathcal{B}_j), \\ d_{P_u(\mathbb{A})}(\mathcal{A}_i, \mathcal{A}_j) &\ge d_{P_u(\mathbb{B})}(\mathcal{B}_i, \mathcal{B}_j). \end{aligned}$$

When studying Δ_ℓ and Δ_u we want to avoid trivial cases where they are infinite because there is a structural obstruction for connectivity given by κ. If κ is reducible then the whole graph $G(n, \kappa, p)$ is disconnected since almost surely there are no edges between the sets $V(\mathcal{A})$ and $V(\mathcal{S} \setminus \mathcal{A})$. Since we want to work with connected graphs, we restrict our attention to the irreducible case.

3. Upper and lower diameters

In this section we study the behavior of the diameters Δ_u and Δ_ℓ. The goal is to prove Theorem 1. We split the proof into two lemmas. We suppose in the two lemmas and the claim of this section that regularity conditions (R) hold.

Given two partitions $\mathbb{A}$ and $\mathbb{B}$, we define their *common refinement* as

$$\mathbb{A} \vee \mathbb{B} := \{\mathcal{A} \cap \mathcal{B} : \mu(\mathcal{A} \cap \mathcal{B}) > 0, \text{ for } \mathcal{A} \in \mathbb{A} \text{ and } \mathcal{B} \in \mathbb{B}\}.$$

(Recall that since we consider essential partitions only, $\mathbb{A} \vee \mathbb{B}$ still remains an essential partition of $\mathcal{S}$.)

Lemma 2. *If $\Delta_\ell < \infty$, then $\Delta_u \leq \Delta_\ell$.*

Proof. Since $\Delta_\ell < \infty$, there exists a partition $\mathbb{A}$ such that $\operatorname{diam} P_\ell(\mathbb{A}) = \Delta_\ell$. Let $\mathbb{B}$ be an arbitrary partition. Consider the common refinement $\mathbb{A} \vee \mathbb{B}$. Since $P_\ell(\mathbb{A} \vee \mathbb{B})$ is a subgraph of $P_u(\mathbb{A} \vee \mathbb{B})$ we have that

$$\operatorname{diam} P_u(\mathbb{A} \vee \mathbb{B}) \leq \operatorname{diam} P_\ell(\mathbb{A} \vee \mathbb{B}).$$

Moreover, we also have that

$$\operatorname{diam} P_u(\mathbb{B}) \leq \operatorname{diam} P_u(\mathbb{A} \vee \mathbb{B}) \qquad \text{and} \qquad \operatorname{diam} P_\ell(\mathbb{A} \vee \mathbb{B}) \leq \operatorname{diam} P_\ell(\mathbb{A}),$$

because $\mathbb{A} \vee \mathbb{B}$ is a refinement of both $\mathbb{A}$ and $\mathbb{B}$. Combining these three inequalities we get that $\operatorname{diam} P_u(\mathbb{B}) \leq \operatorname{diam} P_u(\mathbb{A}) \leq \operatorname{diam} P_\ell(\mathbb{A}) = \Delta_\ell < \infty$. Therefore, the inequality also holds after taking the supremum over all partitions. In particular, we can choose $\mathbb{B}$ such that $\operatorname{diam} P_u(\mathbb{B}) = \Delta_u$, and the desired inequality follows. $\square$

In particular, the above bound gives an easy way to determine Δ_u and Δ_ℓ in the case they are equal. It suffices to find a partition $\mathbb{A}$ for which $\operatorname{diam} P_u(\mathbb{A}) = \operatorname{diam} P_\ell(\mathbb{A})$ holds.

We can also show the following bound.

Lemma 3. *If $\Delta_\ell < \infty$, then $\Delta_\ell \leq \Delta_u + 2$.*

Proof. We state the following claim which we prove below.

Claim 1. *Suppose $\Delta_u < \infty$. Given a partition $\mathbb{B}$, let $\mathcal{B}_s, \mathcal{B}_f \in \mathbb{B}$. There exists a refinement $\mathbb{A} \prec \mathbb{B}$ such that there exist $\mathcal{A}_s, \mathcal{A}_f \in \mathbb{A}$ with $\mathcal{A}_s \subset \mathcal{B}_s$, $\mathcal{A}_f \subset \mathcal{B}_f$ such that $d_{P_\ell(\mathbb{A})}(\mathcal{A}_s, \mathcal{A}_f) \leq \Delta_u$.*

Assuming the claim, the lemma follows easily: indeed, start with a partition $\mathbb{B}$ with $\operatorname{diam} P_\ell(\mathbb{B}) < \infty$. Consider all pairs

$$\mathfrak{P} = \{(\mathcal{B}_i, \mathcal{B}_j) \in \mathbb{B} \times \mathbb{B} : d_{P_\ell(\mathbb{B})}(\mathcal{B}_i, \mathcal{B}_j) > \Delta_u + 2\}.$$

If $\mathfrak{P} = \emptyset$ then we are done, as $\Delta_\ell \leq \max_{\mathcal{B}_i, \mathcal{B}_j \in \mathbb{B}} d_{P_\ell(\mathbb{B})}(\mathcal{B}_i, \mathcal{B}_j) \leq \Delta_u + 2$. Therefore suppose $\mathfrak{P} \neq \emptyset$. Since $\operatorname{diam} P_\ell(\mathbb{B}) < \infty$, given $(\mathcal{B}_i, \mathcal{B}_j) \in \mathfrak{P}$ there exist $\mathcal{B}_s$ and $\mathcal{B}_f$ such that $\kappa_\ell(\mathcal{B}_i, \mathcal{B}_s) > 0$ and $\kappa_\ell(\mathcal{B}_f, \mathcal{B}_j) > 0$.

Since $\Delta_\ell < \infty$, by Lemma 2 we have $\Delta_u \leq \Delta_\ell < \infty$. Thus, by Claim 1, there exists $\mathbb{A} \prec \mathbb{B}$ such that $d_{P_\ell(\mathbb{A})}(\mathcal{A}_s, \mathcal{A}_f) \leq \Delta_u$ for some $\mathcal{A}_s, \mathcal{A}_f$ with $\mathcal{A}_s \subset \mathcal{B}_s$ and $\mathcal{A}_f \subset \mathcal{B}_f$. Then, for any $\mathcal{A}_i, \mathcal{A}_j \in \mathbb{A}$ such that $\mathcal{A}_i \subset \mathcal{B}_i$ and $\mathcal{A}_j \subset \mathcal{B}_j$ we have that $\kappa_\ell(\mathcal{A}_i, \mathcal{A}_s) > 0$ and $\kappa_\ell(\mathcal{A}_f, \mathcal{A}_j) > 0$, and therefore $d_{P_\ell(\mathbb{A})}(\mathcal{A}_i, \mathcal{A}_j) \leq \Delta_u + 2$.

We construct such a partition $\mathbb{A}$ for each pair in $\mathfrak{P}$. Since $\mathfrak{P}$ is finite, consider a common refinement $\mathbb{C}$ of all of these partitions. It is clear that $\mathbb{C}$ has $\max_{\mathcal{C}_i, \mathcal{C}_j \in \mathbb{C}} d_{P_\ell(\mathbb{C})}(\mathcal{C}_i, \mathcal{C}_j) \leq \Delta_u + 2$, and since $\Delta_\ell \leq \max_{\mathcal{C}_i, \mathcal{C}_j \in \mathbb{C}} d_{P_\ell(\mathbb{C})}(\mathcal{C}_i, \mathcal{C}_j)$, the lemma follows. $\square$

We first give a high-level strategy of the proof of Claim 1: we first check whether in the partition $\mathbb{B}$ we find $r \le \Delta_u$, $x_1 \in \mathcal{B}_s$, $x_i \in \mathcal{B}_i$ for $i = 2, \dots, r$, $x_{r+1} \in \mathcal{B}_f$ such that $\kappa(x_i, x_{i+1}) \ge \delta$ for some $\delta > 0$ for all $i = 1, \dots, r$ (this corresponds to a path of length $r \le \Delta_u$ from $\mathcal{B}_s$ to $\mathcal{B}_f$ in the upper partition graph). If this is the case, we easily find the partition $\mathbb{A}$ by considering parts that contain small enough balls around the $r+1$ elements x_i. Otherwise, we refine $\mathbb{B}$ by partitioning each minimal length path from $\mathcal{B}_s$ to $\mathcal{B}_f$; we split each vertex $\mathcal{B}_i$ of the path into 3 parts: one part containing those elements $x_i \in \mathcal{B}_i$ for which there exists a sequence of $x_1 \in \mathcal{B}_s$, $x_j \in \mathcal{B}_j$ for $j = 2, \dots, i-1$ with $\kappa(x_j, x_{j+1}) \ge \delta$ for some $\delta > 0$ for all $j = 1, \dots, i-1$ (this corresponds to a minimal length path from $\mathcal{B}_s$ to this part of $\mathcal{B}_i$), one part containing those elements $x_i \in \mathcal{B}_i$ for which there exists a sequence $x_i \in \mathcal{B}_i$, $x_j \in \mathcal{B}_j$ for $j = i+1, \dots, r$ (for some $r \le \Delta_u$), $x_{r+1} \in \mathcal{B}_f$ with $\kappa(x_j, x_{j+1}) \ge \delta$ for some $\delta > 0$ for all $j = i, \dots, r$ (this corresponds to a minimal length path from this part of $\mathcal{B}_i$ to $\mathcal{B}_f$), and one part containing the remainder of $\mathcal{B}_i$. Crucially, in the upper partition graph the distance of the minimal length path then must increase in this refinement. For the common refinement of all partitions corresponding to all minimal length paths, the argument is repeated, and since Δ_u is finite, after finitely many steps we must find the desired path or reach a contradiction. We now give the detailed proof.

Proof of Claim 1. Eliminate from $\mathcal{S}$ all points where κ is not continuous. Note that the removed set has measure zero and does not affect the calculations of essential infima and suprema.

Next, suppose there exists $\delta > 0$ and suppose that in the remaining set there exist $x_1, \dots, x_{r+1}$ with $r \le \Delta_u$ with $x_1 \in \mathcal{B}_1 := \mathcal{B}_s, x_{r+1} \in \mathcal{B}_{r+1} := \mathcal{B}_f$, $x_i \in \mathcal{B}_i$ for $i = 2, \dots, r$ and $\kappa(x_i, x_{i+1}) \ge \delta$ for $i = 1, \dots, r$. By continuity of κ, there exist $\varepsilon_i, \varepsilon_{i+1} > 0$ such that for any y in the ball $B(x_i, \varepsilon_i)$ and z in the ball $B(x_{i+1}, \varepsilon_{i+1})$ we have $\kappa(y, z) \ge \delta/2$. Consider the partition $\mathbb{A} \prec \mathbb{B}$ in the following way: all parts except for $\mathcal{B}_i$ with $i = 1, \dots, r+1$ remain unchanged: for $i = 1, \dots, r+1$, $\mathcal{B}_i$ is split into $\mathcal{A}_i = \mathcal{B}_i \cap B(x_i, \varepsilon_i)$ and $\mathcal{A}_i' = \mathcal{B}_i \setminus B(x_i, \varepsilon_i)$. Since in $\mathbb{A}$ we have $d_{P_\ell(\mathbb{A})}(\mathcal{A}_s, \mathcal{A}_f) \le r \le \Delta_u$, we found the desired partition $\mathbb{A}$.

Otherwise, there exists no such path of length $r \le \Delta_u$. Consider any shortest path $\mathcal{B}_1 := \mathcal{B}_s, \mathcal{B}_2, \dots, \mathcal{B}_{r+1} := \mathcal{B}_f$ of length $r \le \Delta_u$ in $P_u(\mathbb{B})$. For $i = 1, \dots, r+1$ let

$$\mathcal{A}_i^s := \{x \in \mathcal{B}_i : \exists (x_1, \dots, x_i = x) \in (\mathcal{B}_1, \dots, \mathcal{B}_i) \mid \kappa(x_j, x_{j+1}) > 0, j = 1, \dots, i-1\}$$

be the sets of vertices of $\mathcal{B}_i$ to which there is a path that starts at $\mathcal{B}_s$. Similarly, let

$$\mathcal{A}_i^f := \{x \in \mathcal{B}_i : \exists (x_i = x, \dots, x_{r+1}) \in (\mathcal{B}_i, \dots, \mathcal{B}_{r+1}) \mid \kappa(x_j, x_{j+1}) > 0, j = i, \dots, r\}$$

be the sets of vertices of $\mathcal{B}_i$ from which there is a path that finishes at $\mathcal{B}_f$. Note that for all $i = 1, \dots, r+1$, we must have $\mathcal{A}_i^s \cap \mathcal{A}_i^f = \emptyset$, as otherwise we would have a path of length r between $\mathcal{B}_s$ and $\mathcal{B}_f$ in $P_\ell(\mathbb{B})$, and we would be in the first case addressed in this claim. Consider the partition $\mathbb{A} \prec \mathbb{B}$ induced from splitting $\mathcal{B}_i$ into $\mathcal{A}_i^s, \mathcal{A}_i^f$ and $\mathcal{B}_i \setminus (\mathcal{A}_i^s \cup \mathcal{A}_i^f)$. Since some of these sets might be empty we consider the partition obtained after removing sets of measure zero. Note that for the new partition $\mathbb{A}$, the shortest path starting from $\mathcal{A}_1^s$ and ending at $\mathcal{A}_{r+1}^f$ and using only elements $\mathcal{A}_i^s, \mathcal{A}_i^f, \mathcal{B}_i \setminus (\mathcal{A}_i^s \cup \mathcal{A}_i^f)$ for some $i = 1, \dots, r+1$ must have length $d_1 \ge r+1$ in the upper partition graph corresponding to $\mathbb{A}_1 := \mathbb{A}$ (indeed, if there exists i, and $x \in \mathcal{A}_i^s$, $y \in \mathcal{A}_{i+1}^f$ with $\kappa(x, y) > 0$, then there would exist $\delta > 0$

and $x_1 \in \mathcal{B}_1, \ldots, x_{r+1} \in \mathcal{B}_{r+1}$ with $\kappa(x_i, x_{i+1}) \geq \delta$ for $i = 1, \ldots, r$, and we would be in the first case). If there are several shortest paths of length r in $P_u(\mathbb{B})$ between $\mathcal{B}_s$ and $\mathcal{B}_f$, do independently the same refinement and obtain for each such path a refined partition $\mathbb{A}_i \prec \mathbb{B}$. Note that there are only finitely many partitions, since there are only finitely many shortest paths of length r in $P_u(\mathbb{B})$. As before, when refining, distances in the upper partition graph either stay the same or increase. We may thus take a partition $\mathbb{C}$ which is a common refinement of all $\mathbb{A}_i$, and we have for all $\mathcal{C}_s, \mathcal{C}_f \in \mathbb{C}$ with $\mathcal{C}_s \subset \mathcal{B}_s$ and $\mathcal{C}_f \subset \mathcal{B}_f$, $d_{P_u(\mathbb{C})}(\mathcal{C}_s, \mathcal{C}_f) \geq d_1$. If $d_1 > \Delta_u$, we found a new partition $\mathbb{C}$ with diameter bigger than Δ_u, contradicting the definition of Δ_u. Otherwise we apply the claim with partition $\mathbb{C}$ playing the role of partition $\mathbb{B}$. Note that there are only finitely many elements $\mathcal{C}_s, \mathcal{C}_f$ with $\mathcal{C}_s \subset \mathcal{B}_s$ and $\mathcal{C}_f \subset \mathcal{B}_f$. For a fixed pair of such elements $\mathcal{C}_s, \mathcal{C}_f$ repeat the argument of the proof of the claim (yielding a sequence of refined partitions corresponding to all shortest paths of length d_1 between them), giving either the desired path or a partition being a common refinement of all these partitions. Taking then again the refinement of all refined partitions corresponding to all pairs $\mathcal{C}_s, \mathcal{C}_f$ yields a new refinement $\mathbb{D}$ in which all pairs are at distance $d_2 \geq d_1 + 1$, and the claim can then be applied with $\mathbb{D}$ playing the role of $\mathbb{C}$. Since $\Delta_u < \infty$, after finitely many iterations we must have found the desired path of length at most Δ_u, and the claim follows. □

The first part of Theorem 1 follows now easily by combining Lemma 2 and Lemma 3. For the second part, to show that both bounds can be attained, on the one hand consider a constant kernel defined as $\kappa(x, y) = 1$ for all x, y. Clearly, for any partition $\mathbb{A}$ of $\mathcal{S}$, the upper and lower partitions corresponding to $\mathbb{A}$ are the same graphs, and therefore for such a kernel $\Delta_\ell = \Delta_u$. On the other hand, to show that $\Delta_\ell = \Delta_u + 2$ can be attained, consider the following example: fix $k \geq 2$ and let

$$\mathcal{S} = [0, 1/2] \cup \bigcup_{i=1}^{k+1} \{i\} \cup [k + 3/2, k + 2]$$

$$\mu = \frac{1}{k+3} \left(2\lambda_{[0,1/2]} + \sum_{i=1}^{k+1} \delta_i + 2\lambda_{[k+3/2,k+2]} \right)$$

where δ_i is the Dirac measure and λ_I is the Lebesgue measure restricted to the interval I (by suitable rescaling of intervals we could clearly have the same example in $[0, 1]$); μ is then the new measure obtained from the combinations of Dirac and Lebesgue measures. Define κ for $x \leq y$ as follows (and extend by symmetry):

$$\kappa(x, y) = \begin{cases} 1 & \text{if } y = x + 1 \text{ and } x \in \{1, \ldots, k\}, \\ 1 & \text{if } x \in [0, 1/2] \text{ and } y = 1, \\ x & \text{if } x \in [0, 1/2] \text{ and } y = 2, \\ k + 2 - x & \text{if } x = k \text{ and } y \in [k + 3/2, k + 2], \\ 1 & \text{if } x = k + 1 \text{ and } y \in [k + 3/2, k + 2], \\ 0 & \text{otherwise.} \end{cases}$$

It is easy to see that $\Delta_\ell = k + 2$ and $\Delta_u = k$: indeed, for any partition $\mathbb{A}$, since $\mathbb{A}$ is finite, there must exist an element $\mathcal{A}_0$ of $\mathbb{A}$ such that $\mu(\mathcal{A}_0 \cap [0, \varepsilon]) > 0$ for all $\varepsilon > 0$. Therefore, $\operatorname{ess\,inf}_{x \in \mathcal{A}_0} \kappa(x, 2) = 0$. Analogously, there must also exist an element $\mathcal{A}_{k+2}$ of $\mathbb{A}$ such that $\mu(\mathcal{A}_{k+2} \cap [k + 2 - \varepsilon, k + 2]) > 0$ for all $\varepsilon > 0$. Therefore,

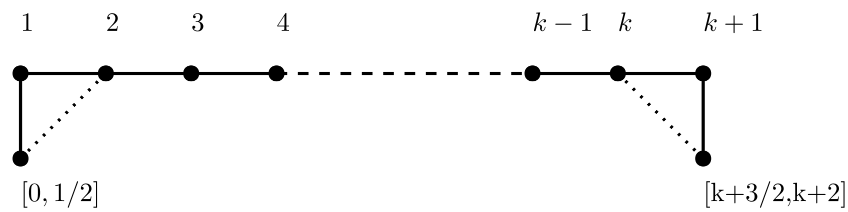

FIGURE 1. The partition graphs $P_\ell(\mathbb{A})$ and $P_u(\mathbb{A})$ only differ in the dotted edges at the endpoints.

$\operatorname{ess\,inf}_{x \in \mathcal{A}_{k+2}} \kappa(x,k) = 0$. This implies that there cannot be a path of length strictly less than $k+2$ between $\mathcal{A}_0$ and $\mathcal{A}_{k+2}$ in $P_\ell(\mathbb{A})$, and hence $\Delta_\ell \geq k+2$. To show that $\Delta_u \leq k$, consider any partition $\mathbb{A}$ and consider two parts $\mathcal{A}_i, \mathcal{A}_j$ such that there exist $\varepsilon > 0$ with $\mu(\mathcal{A}_i \cap [0,\varepsilon]) = \mu(\mathcal{A}_j \cap [0,\varepsilon]) = \mu(\mathcal{A}_i \cap [k+2-\varepsilon, k+2]) = \mu(\mathcal{A}_j \cap [k+2-\varepsilon, k+2]) = 0$. Then, clearly there exists a path of length at most k in $P_u(\mathbb{A})$ between $\mathcal{A}_i$ and $\mathcal{A}_j$. Next, suppose that $\mathcal{A}_0$ is such that $\mu(\mathcal{A}_0 \cap [0,\varepsilon]) > 0$ for all $\varepsilon > 0$. Then, $\operatorname{ess\,sup}_{x \in \mathcal{A}_0} \kappa(x,2) > 0$. Similarly, suppose that $\mathcal{A}_{k+2}$ is such that $\mu(\mathcal{A}_{k+2} \cap [k+2-\varepsilon, k]) > 0$ for all $\varepsilon > 0$. Then, $\operatorname{ess\,sup}_{x \in \mathcal{A}_{k+2}} \kappa(x,k) > 0$. Therefore, we can connect $\mathcal{A}_0(\mathcal{A}_{k+2}$, respectively) with any other element of $\mathbb{A}$ in $P_u(\mathbb{A})$ in at most k steps, and thus $\Delta_u \leq k$. In fact, $\Delta_u = k$ and $\Delta_\ell = k+2$ by taking the partition $\mathbb{A} = \{[0,1/2], \{1\}, \dots, \{k+1\}, [k+3/2, k+2]\}$.

The proof of Theorem 1 is complete.

4. EXPANSION OF NEIGHBORHOODS

In this section we prove that the number of vertices at distance k from a given vertex grows exponentially with k, showing that there exist pairs of vertices that are not too close. The main tool we use is a guided exploration process together with concentration inequalities. In all lemmas of this section we assume as before that regularity conditions (R) hold.

Definition 2. Given a partition $\mathbb{A}$, a configuration of points $X_1, \dots, X_n$ is called $\mathbb{A}$-*balanced* if for each part $\mathcal{A} \in \mathbb{A}$ we have $|V(\mathcal{A})| \geq \mu(\mathcal{A})n/2$.

Lemma 4. *A randomly chosen configuration $X_1, \dots, X_n$ is $\mathbb{A}$-balanced with probability $1 - e^{-\Omega(n)}$.*

Proof. Consider a fixed part $\mathcal{A}$ of $\mathbb{A}$. Every element X_i belongs to $\mathcal{A}$ with probability $\mu(\mathcal{A})$, and therefore, the expected number of X_i in $\mu(\mathcal{A})$ is $n\mu(\mathcal{A})$. Moreover, all elements are independent, and the result follows by Chernoff bounds together with a union bound over all parts $\mathcal{A}$. □

Let $\chi = \{X_1, \dots, X_n\}$ be a randomly chosen configuration in $\mathcal{S}^n$ (that is, each X_i is chosen independently according to μ). We write $\mathbf{P}_\chi(\cdot) = \mathbf{P}(\cdot \mid \chi)$ as shorthand notation. We need a few definitions: let

$$\kappa_{\min} := \min\{\kappa_\ell(\mathcal{A}_i, \mathcal{A}_j) : (\mathcal{A}_i, \mathcal{A}_j) \text{ is an edge of } P_\ell(\mathbb{A})\}$$

and for a partition $\mathbb{A} = \{\mathcal{A}_1, \dots, \mathcal{A}_r\}$ set

$$\|\mathbb{A}\|_\mu = \min\{\mu(\mathcal{A}_i) : \mathcal{A}_i \in \mathbb{A}\}.$$

Definition 3. For a given a partition $\mathbb{A}$, we define the *critical length* as

$$L_c = \max\{t : \kappa_{\min} p(\kappa_{\min}\|\mathbb{A}\|_\mu np/16)^t \le 1\}.$$

Also, a sequence $\mathcal{A}_0, \mathcal{A}_1, \dots, \mathcal{A}_q$ is an *expansion walk* in $P_\ell(\mathbb{A})$ if $\kappa_\ell(\mathcal{A}_t, \mathcal{A}_{t+1}) > 0$ for all t, and the sequence does not repeat parts after the L_c-th step, that is, $\mathcal{A}_s \neq \mathcal{A}_t$ for $s \neq t$ and $s > L_c$.

Remark 3. *Note the following bounds on L_c.*

(i) $\Phi - 3 \le L_c$.
By the second assumption of (1), $(np)^{\Phi-1}/n \le (\log n)/\omega$. *Also* $\kappa \le 1$, *thus*

$$\kappa_{\min} p(\kappa_{\min}\|\mathbb{A}\|_\mu np/16)^{\Phi-3} \le \kappa_{\min} p(\kappa_{\min}\|\mathbb{A}\|_\mu/16)^{\Phi-3} \frac{n\log n}{\omega(np)^2} \le \frac{\log n}{\omega(np)} \le 1,$$

where the last inequality follows from our assumption of $np \ge \omega \log^3 n$.

(ii) If $e^{\log^2(np)} \ge \omega n$ *then* $L_c \le \Phi - 1$.
By the first assumption of (1), $(np)^\Phi \ge \omega n \log n$, *and therefore*

$$\kappa_{\min} p(\kappa_{\min}\|\mathbb{A}\|_\mu np/16)^\Phi = \Omega\left(\omega np \log n (\kappa_{\min}\|\mathbb{A}\|_\mu/16)^\Phi\right) > 1,$$

where the last inequality follows from the fact that $\Phi = \Theta(\log n/\log(np))$ *and the assumption* $e^{\log^2(np)} \ge \omega n$.

(iii) If $e^{\log^2(np)} \le \omega n$ *then* $L_c \le \Phi(1+\alpha(n))$ *with* $\alpha(n) = o(1) \cap \omega(1/\log(np))$.
By the same argument, for such a function $\alpha(n)$,

$$\begin{aligned}\kappa_{\min} p(\kappa_{\min}\|\mathbb{A}\|_\mu np/16)^{\Phi(1+\alpha(n))} &= \Omega\left(\omega (np)^{\Phi\cdot\alpha(n)} \log n (\kappa_{\min}\|\mathbb{A}\|_\mu/16)^{\Phi(1+\alpha(n))}\right) \\ &= \Omega\left(e^{\Phi\cdot\alpha(n)\log(np)} e^{-\Theta(\log n/\log(np))}\right) > 1.\end{aligned}$$

Given an expansion walk, we define a guided exploration process on the set of vertices. At each step t we maintain a partition of the vertex set into three subsets Γ_t, E_t, N_t (active, explored, and neutral): explored vertices and edges incident to them are discarded in step t (and from step t on), edges between active and neutral vertices have not been exposed yet, and edges between two neutral vertices have not been exposed yet either. In step t edges between active vertices and (some) neutral vertices are exposed in a fixed order. Our guided exploration process has two phases. While $t \le L_c + 1$, we obtain a set $\Gamma_t \subset V(\mathcal{A}_t)$ such that $|\Gamma_t| = (\kappa_{\min}\|\mathbb{A}\|_\mu np/16)^t$ (we assume for simplicity that $\kappa_{\min}\|\mathbb{A}\|_\mu np/16$ is an integer). In this phase, we expand the size of the previous subset by a factor of $\kappa_{\min}\|\mathbb{A}\|_\mu np/16$ at each step. In the second case, for $t > L_c + 1$, we obtain subsets of linear size and we visit a new part at each step.

More precisely, let $u \in V(\mathcal{A}_0)$. Define $\Gamma_0 = \{u\}$, $E_0 = E_1 = \emptyset$, and $N_0 = V \setminus \{u\}$. To define Γ_t recursively for $t \ge 1$, we consider the vertices in $V(\mathcal{A}_t) \setminus E_t$ in order until we reach a certain size. Denote the desired size at step t by

$$s(t) = \begin{cases} (\kappa_{\min}\|\mathbb{A}\|_\mu np/16)^t, & \text{if } t \le L_c + 1, \\ \|\mathbb{A}\|_\mu n/16, & \text{otherwise.} \end{cases}$$

(We suppose for simplicity that $s(t)$ is an integer for all t.) Define $\Gamma_t^0 = \emptyset$ and let the event $\mathcal{F}_0 = \emptyset$. For $s \le s(t)$, let

$$u_t^s = \min\{v \in V(\mathcal{A}_t) \setminus (E_t \cup \Gamma_t^{s-1}) : d_G(v, \Gamma_{t-1}) = 1\},$$

if such an element exists, and let $\Gamma_t^s = \Gamma_t^{s-1} \cup \{u_t^s\}$. Otherwise, we fail at step t while trying to add the s-th element and we stop the process. Denote this event by $\mathcal{F}_{t,s}$. Define the event of failing by step t as

$$\mathcal{F}_t = \bigcup_{j=1}^{t-1} \mathcal{F}_j \cup \bigcup_{s=1}^{s(t)} \mathcal{F}_{t,s}.$$

Finally, if $\mathcal{F}_t$ does not hold, let

$$\begin{aligned} \Gamma_t &= \Gamma_t^{s(t)} \setminus E_t, \\ N_t &= N_{t-1} \setminus (\Gamma_t \cup E_t), \\ E_{t+1} &= E_t \cup \Gamma_t. \end{aligned}$$

(Recall that, for $t > L_c$, we only visit parts that have never been visited before, that is, for $t > L_c$, $\mathcal{A}_t \neq \mathcal{A}_{t'}$ for any $t' \neq t$, and thus all vertices in $V(\mathcal{A}_t)$ are neutral.)

Remark 4. *Note that since there are only finitely many elements in each partition, for every expansion walk of length q, we have $q - L_c \leq L$ for some constant $L > 0$. Note also that if $\mathcal{F}_t$ does not hold, then $|\Gamma_t| = s(t)$.*

In the following lemma we show that if in one step of the expansion walk the set of neutral vertices (restricted to a certain part in the expansion walk) is still large enough, we expand well from an active set (by definition of the exploration process, the active set is always restricted to one part in the expansion walk).

Lemma 5. *Let $\mathbb{A}$ be a partition of $\mathcal{S}$ and χ be $\mathbb{A}$-balanced. Fix an expansion walk $\mathcal{A}_0, \mathcal{A}_1, \ldots, \mathcal{A}_q$ in $P_\ell(\mathbb{A})$. Let $u \in V(\mathcal{A}_0)$ and $t \in \mathbb{N}$. Consider the t-th step of the exploration process. Then,*

$$\mathbf{P}_\chi \left(\mathcal{F}_t \mid \mathcal{F}_{t-1}^c \right) \leq e^{-s(t)/2}.$$

Proof. Since χ is $\mathbb{A}$-balanced, in particular we have $|V(\mathcal{A}_t)| \geq \mu(\mathcal{A}_t)n/2$. Moreover, for $t > L_c$, by definition of an expansion walk, $|N_t \cap V(\mathcal{A}_t)| = |V(\mathcal{A}_t)|$, and for $t \leq L_c$, $|N_t \cap V(\mathcal{A}_t)| \geq \mu(\mathcal{A}_t)n/4 \geq \|\mathbb{A}\|_\mu n/4$, since by definition of L_c and our assumption on $np \geq \omega \log^3 n$,

$$\sum_{t \leq L_c} |\Gamma_t| \leq \sum_{s=0}^{L_c} \frac{1}{\kappa_{\min} p (\kappa_{\min} \|\mathbb{A}\|_\mu np/16)^s} \leq \frac{2}{\kappa_{\min} p} \leq \frac{n}{\log^3 n} \leq \frac{\|\mathbb{A}\|_\mu n}{4}.$$

For $t > L_c + 2$, each of the vertices in $|N_{t-1} \cap V(\mathcal{A}_t)|$ is connected to at least one vertex in Γ_{t-1} with probability at least $1 - (1 - \kappa_\ell(\mathcal{A}_{t-1}, \mathcal{A}_t)p)^{|\Gamma_{t-1}|}$, and this lower bound holds independently for all vertices. For $t \leq L_c + 2$, note that for every vertex $v \in N_{t-1}$, for all but at most a $o(1)$-fraction of the pairs (v, w) with $w \in \Gamma_{t-1}$ there is no knowledge about the presence of the edge vw (we say, the pair has not been *exposed* yet): indeed, since for all $t \leq L_c + 2$, $|\Gamma_{t-1}| = o(n)$ and $np \geq \omega \log^3 n$, all but a $o(1)$-fraction of the vertices in $|\Gamma_{t-1}|$ have not been active at time $t - 2$, and the vertex pairs incident to them have not been exposed yet. Hence, each such vertex is also connected to at least one vertex in Γ_{t-1} with probability at least $1-(1-\kappa_\ell(\mathcal{A}_{t-1},\mathcal{A}_t)p)^{|\Gamma_{t-1}|(1-o(1))}$. Note that

$$1-(1-\kappa_\ell(\mathcal{A}_{t-1},\mathcal{A}_t)p)^{|\Gamma_{t-1}|} \geq 1-(1-\kappa_{\min}p)^{|\Gamma_{t-1}|}.$$

In the case when $t \leq L_c+1$, we have

$$\begin{aligned}\kappa_{\min}p|\Gamma_{t-1}| &= \kappa_{\min}ps(t-1)\\ &= \kappa_{\min}p(\kappa_{\min}\|\mathbb{A}\|_\mu np/16)^{t-1}\\ &\leq \kappa_{\min}p(\kappa_{\min}\|\mathbb{A}\|_\mu np/16)^{L_c} \leq 1,\end{aligned}$$

by definition of L_c. Thus, we have

$$1-(1-\kappa_{\min}p)^{|\Gamma_{t-1}|(1-o(1))} \geq \kappa_{\min}p|\Gamma_{t-1}|/2.$$

The number of neighbors of Γ_{t-1} in $N_{t-1}\cap V(\mathcal{A}_t)$ stochastically dominates a binomial random variable with parameters $\|\mathbb{A}\|_\mu n/4$ and $\kappa_{\min}p|\Gamma_{t-1}|/2$. The desired concentration then holds by applying Chernoff bounds for binomial random variables.

In the case when $t > L_c+1$, we have

$$\kappa_{\min}p|\Gamma_{t-1}| = \kappa_{\min}ps(t-1) \geq \kappa_{\min}p(\kappa_{\min}\|\mathbb{A}\|_\mu np/16)^{L_c+1} > 1$$

by definition of L_c. Hence, in this case,

$$\begin{aligned}1-(1-\kappa_\ell(\mathcal{A}_{t-1},\mathcal{A}_t)p)^{|\Gamma_{t-1}|(1-o(1))} &\geq 1-e^{-\kappa_\ell(\mathcal{A}_{t-1},\mathcal{A}_t)p|\Gamma_{t-1}|(1-o(1))}\\ &\geq 1-e^{-\kappa_{\min}p|\Gamma_{t-1}|(1-o(1))}\\ &\geq 1-e^{-1}\\ &\geq 1/2.\end{aligned}$$

The number of neighbors of Γ_{t-1} in $N_{t-1}\cap V(\mathcal{A}_t)$ stochastically dominates a binomial random variable with parameters $\|\mathbb{A}\|_\mu n/4$ and $1/2$, and again the desired concentration then holds by applying Chernoff bounds for binomial random variables. $\square$

We immediately obtain the following lemma.

Lemma 6. *Let $\mathbb{A}$ be a partition of $\mathcal{S}$. Fix an expansion walk $\mathcal{A}_0,\mathcal{A}_1,\dots,\mathcal{A}_q$ in $P_\ell(\mathbb{A})$. Let $u \in V(\mathcal{A}_0)$. Then*

$$\mathbf{P}\left(|\Gamma_q| < s(q)\right) \leq e^{-\log^3 n}.$$

Proof. First assume that χ is $\mathbb{A}$-balanced. By Lemma 5 we have that

$$\begin{aligned}\mathbf{P}_\chi\left(\mathcal{F}_q^c\right) &\leq \sum_{t=1}^{q}\mathbf{P}_\chi\left(\mathcal{F}_t^c|\mathcal{F}_{t-1}\right)\\ &\leq \sum_{t=1}^{q} e^{-s(t)/2}\\ &\leq \sum_{t=1}^{L_c+1} e^{-(\kappa_{\min}\|\mathbb{A}\|_\mu np/16)^t/2} + \sum_{t=L_c+2}^{q} e^{-\|\mathbb{A}\|_\mu n/32}\\ &= e^{-\Omega(np)} + (q-L_c-1)e^{-\Omega(n)} \leq (1/2)e^{-\log^3 n},\end{aligned}$$

where the last inequality follows from the fact that $q - L_c \leq L$ for some constant L (see Remark 4), and due to the assumption of $np \geq \omega \log^3 n$. Now, denote by $\mathcal{B}$ the event that χ is $\mathbb{A}$-balanced. We have

$$\mathbf{P}\left(|\Gamma_q| < s(q)\right) = \mathbf{E}\left(\mathbf{P}_\chi\left(|\Gamma_q| < s(q)\right)(\mathbf{1}_{\mathcal{B}} + \mathbf{1}_{\mathcal{B}^c})\right),$$

where the expectation is taken with respect to the configuration χ. By the first part of this lemma and by Lemma 4,

$$\begin{aligned}\mathbf{P}\left(|\Gamma_q| < s(q)\right) &= \mathbf{E}\left(\mathbf{P}_\chi\left(|\Gamma_q| < s(q)\right)\mathbf{1}_{\mathcal{B}}\right) + \mathbf{E}\left(\mathbf{P}_\chi\left(|\Gamma_q| < s(q)\right)\mathbf{1}_{\mathcal{B}^c}\right)\\ &\leq (1/2)e^{-\log^3 n} + \mathbf{P}\left(\mathcal{B}^c\right)\\ &\leq (1/2)e^{-\log^3 n} + e^{-\Omega(n)}\\ &\leq e^{-\log^3 n}.\end{aligned}$$

□

Lemma 7. *Let $\mathbb{A}$ be a partition such that $\mathcal{A}_r, \mathcal{A}_s, \mathcal{A}_j$ is an expansion walk in $P_\ell(\mathbb{A})$ and there is an expansion walk of length q from $\mathcal{A}_i$ to $\mathcal{A}_r$. Then, there exists a constant $c > 0$ such that for $u \in V(\mathcal{A}_i), w \in V(\mathcal{A}_j)$,*

$$\mathbf{P}\left(d_G(u,w) > q+2\right) \leq e^{-\left(\min\{(cnp)^{q+2}/n, \log^3 n\}\right)}.$$

Proof. Let $\Gamma(u)$ and $\Gamma(w)$ be the associated guided exploration processes starting from u and w, respectively. Note that by Lemma 6, with probability at least $1 - e^{-\log^3 n}$,

$$|\Gamma_q(u)| = s(q).$$

Also, with probability at least $1 - e^{-\log^3 n}$, $|\Gamma_1(w)| = \|\mathbb{A}\|_\mu \kappa_{\min} np/16$. Assume that both $|\Gamma_q(u)| = s(q)$ and $|\Gamma_1(w)| = \|\mathbb{A}\|_\mu \kappa_{\min} np/16$. For $u = w$ the statement is trivial, thus assume $u \neq w$. If $w \in \Gamma_t(u)$ for some $t \leq q$, then the statement trivially holds as well, so assume also this is not the case.

Case 1: $\Gamma_{q-1}(u) \subseteq \mathcal{A}_s$ and $|\Gamma_{q-1}(u)| = \Omega(|\mathcal{A}_s|/\log^3 n)$. In this case, by our assumption of $np \geq \omega \log^3 n$, with probability at least $1 - e^{-\log^3 n}$, $d_G(u,w) \leq q$: indeed, the probability that there is no edge between $\Gamma_{q-1}(u)$ and $\Gamma_1(w)$ is at most $(1-\kappa_{\min} p)^{\Omega(n^2 p/\log^3 n)} \leq e^{-\Omega((np)^2/\log^3 n)} \leq e^{-\log^3 n}$, and the statement follows.

Case 2: $\Gamma_{q-1}(u) \subseteq \mathcal{A}_s$ and $|\Gamma_{q-1}(u)| = o(|\mathcal{A}_s|/\log^3 n)$. In this case (or also, if in previous steps of the expansion walk starting from u some edges between vertices in $\mathcal{A}_s$ and $\mathcal{A}_r$ were exposed), $\Gamma_{q-1}(u)$ is a randomly chosen subset of $\mathcal{A}_s$, independently of $\Gamma_1(w)$, since by assumption $w \notin \Gamma_t(u)$. Then, by our assumption of $np \geq \omega \log^3 n$, with probability at least $1 - e^{-\log^3 n}$, $|\Gamma_{q-1}(u) \cap \Gamma_1(w)| = o(|\Gamma_1(w)|)$, and at most a $o(1)$-fraction of pairs of vertices in $\Gamma_q(u)$ and $\Gamma_1(w)$ has been exposed. (The same clearly holds if no edge between vertices of $\mathcal{A}_s$ and $\mathcal{A}_r$ has been exposed in the expansion walk starting from u.) Thus, the probability that there is no edge between $\Gamma_q(u)$ and $\Gamma_1(w)$ is at most

$$(1-\kappa_{\min} p)^{(1+o(1))s(q)\|\mathbb{A}\|_\mu \kappa_{\min} np/16} \leq e^{-\left(\min\{(cnp)^{q+2}/n, c(np)^2\}\right)}$$

for a sufficiently small constant $c > 0$.

In all cases, noting that $(np)^2 \geq \omega \log^3 n$, the lemma follows by summing the failure probabilities. □

5. Bounding the diameter

We dedicate this section to the proof of Theorem 2.

We break up the proof of the theorem into seven lemmas where we study the behavior of the diameter of $G(n,\kappa,p)$ depending on the relationships between Φ, Δ_u and Δ_ℓ. Once more, we assume in all lemmas that regularity conditions (R) hold.

Lemma 8. *With high probability,* $\operatorname{diam} G(n,\kappa,p) \geq \Delta_u$.

Proof. Consider a partition $\mathbb{A}$ attaining Δ_u. With probability 1, there is no edge between any pair of vertices w, x and any two elements $\mathcal{A}_k, \mathcal{A}_\ell \in \mathbb{A}$, such that $w \in V(\mathcal{A}_k), x \in V(\mathcal{A}_\ell)$, and $\mathcal{A}_k, \mathcal{A}_\ell \notin E(P_u(\mathbb{A}))$. Note that with probability $1 - e^{-\Omega(n)}$ we can find two vertices $u, v \in V(\mathcal{A}_i), V(\mathcal{A}_j)$ such that $\mathcal{A}_i$ and $\mathcal{A}_j$ are at distance Δ_u in $P_u(\mathbb{A})$, and hence the lemma follows. $\square$

Lemma 9. *Suppose $\Delta_\ell > \Phi$. With high probability,* $\operatorname{diam} G(n,\kappa,p) \leq \Delta_\ell$.

Proof. We may assume $\Delta_\ell < \infty$, as otherwise there is nothing to prove. Consider a partition $\mathbb{A}$ attaining Δ_ℓ, and consider two arbitrary vertices $u \in V(\mathcal{A}_i)$ and $w \in V(\mathcal{A}_j)$. There exists $\mathcal{A}_r \in \mathbb{A}$ such that $\mathcal{A}_r, \mathcal{A}_s, \mathcal{A}_j$ is an expansion walk of length 2 in $P_\ell(\mathbb{A})$, and such that there exists an expansion walk of length q from $\mathcal{A}_i$ to $\mathcal{A}_r$, for some $\Phi - 2 \leq q \leq \Delta_\ell - 2$: we may assume the upper bound, since if no such walk of length at most $\Delta_\ell - 2$ would exist, then also no path of length at most $\Delta_\ell - 2$ between $\mathcal{A}_i$ and $\mathcal{A}_r$ would exist, and then $\mathcal{A}_i$ and $\mathcal{A}_j$ would be at distance bigger than Δ_ℓ, contradicting the fact that the maximal distance between $\mathcal{A}_i$ and $\mathcal{A}_j$ in $P_\ell(\mathbb{A})$ is at most Δ_ℓ. On the other hand, we clearly may assume the lower bound, since for $q \leq \Phi - 3$ we may reuse partitions to make the walk longer: more precisely, if the expansion walk had length $\Phi - 3$ or less, by making zigzags between the first two partitions of the walk we can make the walk longer and end up with a walk of length $\Phi - 2$ or $\Phi - 1$; indeed, note that $\Delta_\ell \geq 3$, hence there exist at least four partitions, and the last two partitions of the walk can be chosen to appear only once; all other partitions can be reused, since by Remark 3(i), $\Phi - 3 \leq L_c$. By Lemma 7,

$$\mathbf{P}\left(d_G(u,w) > q+2\right) \leq e^{-\left(\min\{(cnp)^{q+2}/n, \log^3 n\}\right)} \leq n^{-\omega},$$

where the second inequality uses the fact that $q + 2 \geq \Phi$ and $(cnp)^\Phi/n \geq \omega \log n$ which follows from the first assumption of (1), noting that $\Phi = O(1)$. By a union bound over all pairs of vertices, the statement follows. $\square$

Lemma 10. *With high probability,* $\operatorname{diam} G(n,\kappa,p) \geq \Phi$.

Proof. Since $\kappa \leq 1$, we can couple $G(n,\kappa,p)$ so that it is a subgraph of $G(n,1,p)$. This can be done, for instance, by using uniform random variables $U_{ij} \in [0,1]$ and letting each edge be present if $U_{ij} < \kappa(X_i, X_j)p$ in $G(n,\kappa,p)$, and if $U_{ij} < p$ in $G(n,1,p)$. Note that $G(n,1,p)$ is nothing but $G(n,p)$. Recall that by Lemma 1 we have $\operatorname{diam} G(n,p) = \Phi$. Hence, with high probability we have $\operatorname{diam} G(n,\kappa,p) \geq \operatorname{diam} G(n,p) = \Phi$. $\square$

Lemma 11. *Suppose $\Phi \geq \Delta_u$. If there exists a partition $\mathbb{A}$ and $i \leq j$, such that there exists no path of length exactly Φ between $\mathcal{A}_i$ and $\mathcal{A}_j$ in $P_u(\mathbb{A})$, then with high probability* $\operatorname{diam} G(n,\kappa,p) \geq \Phi + 1$.

Proof. Fix an arbitrary partition $\mathbb{A}$ and let $i \leq j$ as in the statement of the lemma. First recall that since $\kappa \leq 1$, by the same argument as in Lemma 10, $G(n,\kappa,p)$ can be coupled to be a subgraph of $G(n,p)$. Let Y be the random variable counting the number of vertices $y \in V(\mathcal{A}_i)$ such that there exists $z \in V(\mathcal{A}_j)$ at distance bigger than $\Phi - 1$ in $G(n,\kappa,p)$, and let Z be the same random variable in the coupled $G(n,p)$, so that $Y \geq Z$.

Our proof is based in the second moment method, using ideas similar to Theorem 10.10 of [6]. From now on we consider the coupled $G(n,p)$ graph. First we want to bound from below the first moment of Z. Denoting by $\mathcal{B}$ the event that χ is $\mathbb{A}$-balanced, we have that

$$\mathbf{E}Z \geq \mathbf{E}(Z \mid \mathcal{B})\mathbf{P}\left(\mathcal{B}\right). \tag{2}$$

Note that $\mathbf{P}\left(\mathcal{B}\right) = 1 + o(1)$ since $\mathbf{P}\left(\mathcal{B}^c\right) \leq e^{-\Omega(n)}$. Now we condition on the point configuration χ, which we can assume to be $\mathbb{A}$-balanced. For $y \in V(\mathcal{A}_i)$, let $\mathcal{R}_y$ be the event that there exists $z \in V(\mathcal{A}_j)$ such that $d_G(y,z) \geq \Phi$. We have

$$\mathbf{E}_\chi Z = \sum_{y \in V(\mathcal{A}_i)} \mathbf{P}_\chi\left(\mathcal{R}_y\right).$$

Given a vertex $v \in V(G)$, define

$$B(v) = \{u : d_G(u,v) \leq \Phi - 2\} \quad \text{and} \quad S(v) = \{u : d_G(u,v) = \Phi - 2\}.$$

We say that the event $\mathcal{E}_y$ holds if we have

$$|B(y)| \leq 2(np)^{\Phi-2} \quad \text{and} \quad |S(y)| = (1+o(1))(np)^{\Phi-2}.$$

To bound the probability of $\mathcal{R}_y$ from below, we note that

$$\mathbf{P}_\chi\left(\mathcal{R}_y\right) \geq \mathbf{P}_\chi\left(\mathcal{R}_y \mid \mathcal{E}_y\right)\mathbf{P}_\chi\left(\mathcal{E}_y\right).$$

Conditional under $\mathcal{E}_y$, for a vertex $z \in V(\mathcal{A}_j)$ outside of $B(y)$, the probability that it does not connect to any vertex in $S(y)$ is

$$(1-p)^{(1+o(1))(np)^{\Phi-2}}.$$

Since $(np)^{\Phi-2} = o(n)$ and χ is $\mathbb{A}$-balanced, almost all vertices of $V(\mathcal{A}_j)$ are outside $B(y)$. Moreover, since the edges for different vertices $z \in V(\mathcal{A}_j)$ outside of $B(y)$ are all independent,

$$\mathbf{P}_\chi\left(\mathcal{R}_y^c \mid \mathcal{E}_y\right) = \left(1 - (1-p)^{(1+o(1))(np)^{\Phi-2}}\right)^{|V(\mathcal{A}_j)|(1+o(1))}.$$

From the second assumption in Condition (1) of $(np)^{\Phi-1}/n \leq \log n/\omega$ we have

$$(1-p)^{(1+o(1))(np)^{\Phi-2}} \geq (1-p)^{(1+o(1))\log n/(\omega p)} \geq e^{-c\log n/\omega} \gg 1/\sqrt{n},$$

where the second inequality follows from the fact that $1 - p \geq e^{-cp}$ for c small enough and $p \leq 1/2$. Thus, $1 - (1-p)^{(1+o(1))(np)^{\Phi-2}} \leq e^{-1/\sqrt{n}}$. Then, we have that $\mathbf{P}_\chi\left(\mathcal{R}_y^c \mid \mathcal{E}_y\right) \leq e^{-\Omega(\sqrt{n})}$. Since $\mathbf{P}_\chi\left(\mathcal{E}_y^c\right) \leq e^{-(\log n)^3/\omega}$ for any arbitrarily slowly growing function ω by our assumption that $np \geq \omega \log^3 n$ (the desired concentration for $|B(y)|$ and $|S(y)|$ for a fixed y follows by expanding neighborhoods inductively using a breadth first search, and then Chernoff bound is applied), we have $\mathbf{P}_\chi\left(\mathcal{E}_y\right) = 1 + o(1)$, and thus $\mathbf{P}_\chi\left(\mathcal{R}_y\right) = 1 + o(1)$.

So, if χ is $\mathcal{A}$-balanced, we have that $\mathbf{E}_\chi Z = (1+o(1))|V(\mathcal{A}_i)|$. This implies

$$\mathbf{E}(Z \mid \mathcal{B}) = \mathbf{E}\left(\mathbf{E}_\chi Z \mid \mathcal{B}\right) = (1+o(1))\mathbf{E}|V(\mathcal{A}_i)|.$$

Going back to (2), we have

$$\mathbf{E}Z \geq (1+o(1))\mathbf{E}|V(\mathcal{A}_i)| = (1+o(1))\mu(\mathcal{A}_i)n \to \infty.$$

We bound the second moment trivially from above as follows:

$$\mathbf{E}_\chi Z^2 = \sum_{y \in V(\mathcal{A}_i)} \sum_{y' \in V(\mathcal{A}_i)} \mathbf{P}_\chi \left(\mathcal{R}_y \cap \mathcal{R}_{y'}\right) \leq |V(\mathcal{A}_i)|^2.$$

Therefore, averaging over χ we get

$$\mathbf{E}Z^2 = \mathbf{E}\left(\mathbf{E}_\chi Z^2\right) \leq \mathbf{E}|V(\mathcal{A}_i)|^2 = \mu(\mathcal{A}_i)^2 n^2 + \mu(\mathcal{A}_i)(1-\mu(\mathcal{A}_i))n.$$

Thus, by the Cauchy-Schwartz inequality we have

$$\mathbf{P}\left(Z > 0\right) \geq \frac{(\mathbf{E}Z)^2}{\mathbf{E}Z^2} \geq \frac{(1+o(1))\mu(\mathcal{A}_i)^2 n^2}{\mu(\mathcal{A}_i)^2 n^2 + \mu(\mathcal{A}_i)(1-\mu(\mathcal{A}_i))n} \to 1.$$

By the coupling we are using, $\mathbf{P}\left(Y > 0\right) \to 1$ also. (In fact, $Y = \Omega(n)$ with probability tending to 1.) Therefore, with high probability there exists a remote pair of vertices $u \in V(\mathcal{A}_i)$ and $v \in V(\mathcal{A}_j)$ with $d(u,v) \geq \Phi$. Since there is no path of length exactly Φ between $\mathcal{A}_i$ and $\mathcal{A}_j$ in $P_u(\mathbb{A})$, with high probability, v cannot be reached by u in Φ steps. Hence, the distance between u and v is at least $\Phi+1$, and the lemma follows. □

Lemma 12. *Suppose* $3 \leq \Delta_\ell \leq \Phi = O(1)$. *With high probability,* $\operatorname{diam} G(n,\kappa,p) \leq \Phi + 1$.

Proof. Following the notation and proof of Lemma 9, in this case there exists an expansion walk of length $q = \Phi - 2$ or $q = \Phi - 1$ from $\mathcal{A}_i$ to $\mathcal{A}_r$ (the lower bound is as in Lemma 9. If the upper bound did not hold, then there would be no expansion walk of length at most $\Phi - 1$ between $\mathcal{A}_i$ and $\mathcal{A}_r$, hence also no path of length at most $\Phi - 1$ of length $\mathcal{A}_r$, and then $\mathcal{A}_i$ and $\mathcal{A}_j$ would be at distance bigger than $\Phi + 1 > \Delta_\ell$, contradicting the fact that the maximal distance between $\mathcal{A}_i$ and $\mathcal{A}_j$ in $P_\ell(\mathbb{A})$ is at most Δ_ℓ). Moreover,

$$\mathbf{P}\left(d_G(u,w) > q+2\right) \leq e^{-\left(\min\{(cnp)^{q+2}/n, \log^3 n\}\right)} \leq n^{-\omega},$$

still holds, since $\Phi = O(1)$ and thus $(cnp)^\Phi/n \geq c'\omega' \log n$ for some constants $c, c' > 0$ and some function ω' tending to infinity with n, so that $e^{-c'\omega'\log n} = n^{-\omega}$. The statement follows as before. □

Lemma 13. *Suppose* $\Delta_\ell \leq \Phi = O(1)$. *If there exists a partition* $\mathbb{A}$ *such that for any* $i \leq j$, *there exists a path of length exactly* Φ *between* $\mathcal{A}_i$ *and* $\mathcal{A}_j$ *in* $P_\ell(\mathbb{A})$, *then with high probability* $\operatorname{diam} G(n,\kappa,p) \leq \Phi$.

Proof. Following the notation and proof of Lemma 9, in this case there exists an expansion walk of length exactly $q = \Phi - 2$ from $\mathcal{A}_i$ to $\mathcal{A}_r$. The argument is then as in Lemma 12. □

Lemma 14. *Suppose* $3 \leq \Delta_\ell < \infty$ *and* $\Phi = \omega(1)$. *Then* $\operatorname{diam} G(n,\kappa,p) \leq \Phi(1 + o(1))$ *with high probability.*

Proof. The argument is as in Lemma 9. Recall that $(np)^\Phi/n \geq \omega \log n$ which follows from the first assumption of (1) and recall also that $np \geq \omega \log^3 n$. Let f be a function so that $f = o(1) \cap \omega(1/\log\log n)$. Observe that

$$(cnp)^{\Phi(1+f)}/n \geq c^\Phi (cnp)^{f\Phi} \omega \log n \geq \omega \log n,$$

since $(c(cnp)^f)^\Phi \geq (c(\log n)^{3f})^\Phi$ which follows from the fact that $3f \log\log n = \omega(1)$ and thus $3f \log\log n + \log c > 0$.

Thus, for $q + 2 = \lceil \Phi(1+f) \rceil$ (note that since $\Phi = \Omega(1)$, $q + 2 = \Phi(1 + o(1))$, by Lemma 7,

$$\mathbf{P}\left(d_G(u,w) > q+2\right) \leq e^{-\left(\min\{(cnp)^{q+2}/n, \log^3 n\}\right)} \leq n^{-\omega},$$

and the lemma follows. □

Remark 5. *Note that for $\Phi = \omega(1)$ an upper bound of* $\operatorname{diam} G(n,\kappa,p) \leq \Phi + 1$ *does not hold in general. An inhomogeneous random graph $G(n,\kappa,p)$ with constant kernel $\kappa = c < 1$ has diameter equal to the diameter of the Erdős-Rényi model $G(n,cp)$, and the diameter of $G(n,cp)$ can already be bigger than $\Phi+1$ (for c sufficiently small). For a more precise upper bound of a concrete inhomogeneous random graph we would need information about κ and about $\|\mathbb{A}\|_\mu$.*

Finally, combining all the lemmas above we obtain the proof of the three statements in Theorem 2.

Proof of Theorem 2. By combining Lemma 8 and 9 (using also Lemma 2) (i) follows. The first part of (ii) follows by Lemma 10 and Lemma 9, and the second part follows by adding Lemma 11. The first part of (iii) follows by Lemma 10 and 12; for the second part Lemma 11 and Lemma 13 is used. Finally, (iv) holds by Lemma 10 and Lemma 14. □

Acknowledgements. The authors would like to thank the anonymous referees for their careful reading and very valuable comments, which greatly improved this manuscript. The work of the first author was made possible through generous support by NSERC. Part of this work was done at the Centre de Recerca Matemàtica (CRM) in Barcelona. The authors are very grateful for its hospitality. The second author would also like to thank the Center for Mathematical Modeling (CNRS UMI 2807) at Universidad de Chile. This work was finished during his visit.

References

1. L. Addario-Berry, N. Broutin, and C. Goldschmidt, *Critical random graphs: limiting constructions and distributional properties*, Electronic Journal of Probability **15** (2010), 741–775.
2. ———, *The continuum limit of critical random graphs*, Probability Theory and Related Fields **152** (2012), no. 3-4, 367–406.
3. S. Bhamidi, N. Broutin, S. Sen, and X. Wang, *Scaling limits of random graph models at criticality: Universality and the basin of attraction of the Erdős–Rényi random graph*, Preprint. arXiv:1411.3417 [math.PR].
4. S. Bhamidi, S. Sen, and X. Wang, *Continuum limit of critical inhomogeneous random graphs*, Probability Theory and Related Fields **169** (2017), 565–641.
5. B. Bollobás, *The diameter of random graphs*, Transactions of the American Mathematical Society **267** (1981), no. 1, 41–52.
6. ———, *Random graphs*, Cambridge University Press, 2001.

7. B. Bollobás, S. Janson, and O. Riordan, *The phase transition in inhomogeneous random graphs*, Random Structures & Algorithms **31** (2007), 3–122.
8. C. Borgs, J.T. Chayes, L. Lovász, V.T. Sós, and K. Vesztergombi, *Convergent sequences of dense graphs I: Subgraph frequencies, metric properties and testing*, Advances in Mathematics **219** (2008), no. 6, 1801–1851.
9. ______, *Convergent sequences of dense graphs II: Multiway cuts and statistical physics*, Annals of Mathematics **176** (2012), 151–219.
10. F. Chung and L. Lu, *The diameter of sparse random graphs*, Advances in Applied Mathematics **26** (2001), 257–279.
11. L. Devroye and N. Fraiman, *Connectivity threshold of inhomogeneous random graphs*, Random Structures & Algorithms **45** (2014), no. 3, 408–420.
12. M. Doležal, J. Hladký, and A. Máthé, *Cliques in dense inhomogeneous random graphs*, Random Structures & Algorithms **51** (2017), no. 2, 275–314.
13. V. Klee and D. Larman, *Diameters of random graphs*, Canadian Journal of Mathematics **33** (1981), 618–640.
14. L. Lovász, *Large networks and graph limits*, American Mathematical Society, 2012.
15. L. Lovász and B. Szegedy, *Limits of dense graph sequences*, Journal of Combinatorial Theory, Series B **96** (2006), no. 6, 933–957.
16. A. Nachmias and Y. Peres, *Critical random graphs: Diameter and mixing time*, The Annals of Probability **36** (2008), no. 4, 1267–1286.
17. I. Norros and H. Reittu, *On a conditionally poissonian graph process*, Advances in Applied Probability **38** (2006), no. 1, 59–75.
18. O. Riordan and N. Wormald, *The diameter of sparse random graphs*, Combinatorics, Probability & Computing **19** (2010), no. 5-6, 835–926.
19. B. Söderberg, *General formalism for inhomogeneous random graphs*, Physical review E **66** (2002), no. 6, 066121.

Department of Statistics and Operations Research, University of North Carolina at Chapel Hill, 337 Hanes Hall, Chapel Hill, NC 27599, USA

E-mail address: `fraiman@email.unc.edu`

Laboratoire J–A Dieudonné, Université Nice Sophia Antipolis, Parc Valrose, 06108 Nice Cedex 02, France

E-mail address: `dmitsche@unice.fr`