

THE DIAMETER OF INHOMOGENEOUS RANDOM GRAPHS

NICOLAS FRAIMAN AND DIETER MITSCHKE

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 $\kappa : \mathcal{S} \times \mathcal{S} \rightarrow [0, 1]$ 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 $V = \{1, \dots, n\}$ 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, \dots, 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

$$\begin{aligned}\kappa_\ell(\mathcal{A}, \mathcal{B}) &= \text{ess inf}\{\kappa(x, y) : x \in \mathcal{A}, y \in \mathcal{B}\}, \\ \kappa_u(\mathcal{A}, \mathcal{B}) &= \text{ess sup}\{\kappa(x, y) : x \in \mathcal{A}, y \in \mathcal{B}\}.\end{aligned}$$

For a partition $\mathbb{A} = \{\mathcal{A}_1, \dots, \mathcal{A}_m\}$ 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 the graph with vertex set \mathbb{A} and where $(\mathcal{A}_i, \mathcal{A}_j)$ is an edge if $\kappa_u(\mathcal{A}_i, \mathcal{A}_j) > 0$.

A kernel κ on (\mathcal{S}, μ) 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*.

Date: September 16, 2021.

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 University of Pennsylvania.

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 $n \rightarrow \infty$. Since we are only interested in results that hold with high probability, we avoid working with partitions that have measure zero sets. Therefore, every time we refer to a partition we assume that it has no measure zero sets.

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 $\text{diam } G = \max_{u,v} d_G(u, v)$. We study the diameter of $G(n, \kappa, p)$ by studying the diameters of the induced graphs $P_\ell(\mathbb{A})$ and $P_u(\mathbb{A})$.

Our main result is then stated in terms of the following two constants

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

where \mathbb{A} ranges over all partitions with no measure zero sets. In Section 3 we show our first result:

Theorem 1. *Suppose $\Delta_\ell < \infty$. Then*

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

both bounds being tight.

In order to introduce our next theorem we need to 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 [6]. In order to simplify our main result (see below for more details), we will assume that $\log n / \log np$ is at least $2 \log \log n / \log np$ away from every integer.

Our main result is the following:

Theorem 2. *Let κ be an irreducible kernel that is continuous $(\mu \otimes \mu)$ -almost everywhere and let $G(n, \kappa, p)$ be the induced inhomogeneous random graph. With high probability, the following statements hold:*

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

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

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

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

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

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

(iii) *If $\Delta_\ell \leq \Phi$, then*

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

Moreover,

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

iff for every partition \mathbb{A} there exist \mathcal{A}_i and \mathcal{A}_j , with no walk of length exactly Φ between them in $P_u(\mathbb{A})$.

1.1. Background and history. A discrete version of this model was introduced by Söderberg [20]. The sparse case (when the number of edges is linear in the number n of vertices) was studied in detail in by Bollobás, Janson and Riordan [8]. 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 [13]. 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 [16] and further studied in depth by Borgs, Chayes, Lovász, Sós and Vesztergombi [9, 10] among others. For a thorough introduction to the subject of graph limits see the book by Lovász [15].

The diameter of random graphs has been studied widely. In particular, for the Erdős-Rényi model, Bollobás [6] generalized the results from Klee and Larman [14] characterizing the case of constant diameter. Later, Chung and Lu [12] proved concentration results in various different ranges. More recently, Riordan and Wormald [19] 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 [17] 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 [5] to the Norros-Reittu [18] random graph model, and then further generalized by Bhamidi, Broutin, Sen and Wang [4].

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 [8] with minor changes. We also use the following standard notation: we write $f = O(g)$ if f/g is bounded and $f = o(g)$ if $f/g \rightarrow 0$.

Define $N(u) = \{v \in V : (u, v) \in E\}$ the neighborhood of vertex u , and for a subset $U \subset V$ let $N(U) = \cup_{u \in U} N(u)$. 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.,

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

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 partition $\mathbb{A} = \{\mathcal{A}_1, \dots, \mathcal{A}_m\}$ 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, each $\mathcal{B}_i = \cup_{p=1}^{m_i} \mathcal{A}_p^{(i)}$ 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 exists p, q and $\mathcal{A}_i, \mathcal{A}_j$ 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,

$$\text{diam } P_u(\mathbb{B}) \leq \text{diam } P_u(\mathbb{A}),$$

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,

$$\text{diam } P_\ell(\mathbb{B}) \geq \text{diam } P_\ell(\mathbb{A}).$$

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

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

By the above properties it is clear that the values of $\Delta_\ell = \inf_{\mathbb{A}} \text{diam } P_\ell(\mathbb{A})$ and $\Delta_u = \sup_{\mathbb{A}} \text{diam } P_u(\mathbb{A})$ are well defined. 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. Let

$$\lambda(x) = \int_{\mathcal{S}} \kappa(x, y) d\mu(y) \quad \text{and} \quad \lambda_* = \text{ess inf } \lambda(x).$$

In [13] the constant λ_* is called the *isolation parameter* and it is shown to determine the connectivity threshold. Note that the expected degree of a vertex of type x is $\lambda(x)np$. Our arguments require this value to be $\Theta(np)$ so we will work with kernels that satisfy $\lambda_* > 0$. Observe that the assumption $\lambda_* > 0$ implies that all lower partition graphs $P_\ell(\mathbb{A})$ are connected (the upper partition graphs are connected anyway, since κ is assumed to be irreducible).

Recall that we defined the expansion factor as $\Phi = \lceil \log n / \log np \rceil$. Its relation to the diameter of the Erdős and Rényi model $G(n, p)$ is made precise now. Since we use it in our argument we include the following result obtained first by [6] as a Lemma.

Lemma 1 ([7], Corollary 10.12). *Let $k \geq 2$ be finite. If $(np)^k/n - 2 \log n \rightarrow \infty$ and $(np)^{k-1}/n - 2 \log n \rightarrow -\infty$. Then, with high probability $\text{diam}(G(n, p)) = k$.*

Since we assume that $\log n / \log np$ is at least $2 \log \log n / \log np$ away from every integer, clearly, the diameter of $G(n, p)$ is with high probability concentrated on one value, and hence clearly $k = \Phi$. In fact, this assumption is stronger and implies that

$$(1) \quad \frac{(np)^\Phi}{n} - \omega \log n \rightarrow \infty \quad \text{and} \quad \frac{(np)^{\Phi-1}}{n} - \frac{1}{\omega} \log n \rightarrow -\infty.$$

The 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.

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

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

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 $\text{diam } P_\ell(\mathbb{A}) = \Delta_\ell$. 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

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

Moreover, we also have that

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

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

$$\text{diam } P_u(\mathbb{A}) = \text{diam } P_\ell(\mathbb{A}).$$

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 $c(\mathcal{A}_s) = \mathcal{B}_s$, $c(\mathcal{A}_f) = \mathcal{B}_f$ 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 $\text{diam } P_\ell(\mathbb{B}) < \infty$. 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 $\text{diam } P_\ell(\mathbb{B}) < \infty$ given $(\mathcal{B}_i, \mathcal{B}_j) \in \mathfrak{P}$ there exists \mathcal{B}_s and \mathcal{B}_f such that $\kappa_\ell(\mathcal{B}_i, \mathcal{B}_s) > 0$ and $\kappa_\ell(\mathcal{B}_j, \mathcal{B}_f) > 0$.

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 $c(\mathcal{A}_s) = \mathcal{B}_s$ and $c(\mathcal{A}_f) = \mathcal{B}_f$. Then, for any $\mathcal{A}_i, \mathcal{A}_j \in \mathbb{A}$ such that $c(\mathcal{A}_i) = \mathcal{B}_i$ and $c(\mathcal{A}_j) = \mathcal{B}_j$ 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

Proof of Claim 1. Eliminate from \mathcal{S} all points where κ is not continuous. Note that the removed set has measure 0 and does not affect the calculations of essential infima and suprema.

Next, suppose that in the remaining set there exist x_1, \dots, x_r with $r \leq \Delta_u$ with $x_1 \in \mathcal{B}_1 := \mathcal{B}_s, x_r \in \mathcal{B}_r := \mathcal{B}_f, x_i \in \mathcal{B}_i$ for $i = 2, \dots, r-1$ and $\kappa(x_i, x_{i+1}) > 0$ for $i = 1, \dots, r-1$. 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 $\kappa(y, z) > 0$. Consider the partition $\mathbb{A} \prec \mathbb{B}$ in the following way: all parts except for \mathcal{B}_i with $i = 1, \dots, r$ remain unchanged: for $i = 1, \dots, r$, \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 $d_{P_\ell(\mathbb{A})}(\mathcal{A}_s, \mathcal{A}_f) \leq r \leq \Delta_u$, we found the desired partition \mathbb{A} .

Otherwise, there exists no such path of length $r \leq \Delta_u$. Consider any path of minimal distance $\mathcal{B}_1 := \mathcal{B}_s, \mathcal{B}_2, \dots, \mathcal{B}_r := \mathcal{B}_f$ of length $r \leq \Delta_u$ in $P_u(\mathbb{B})$. For $i = 1, \dots, r$ 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,

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

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 $i = 1, \dots, r$, 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 . 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^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 $i = 1, \dots, r$ must have length $d_1 \geq r+1$ in the upper partition graph corresponding to $\mathbb{A}_1 := \mathbb{A}$. If there are several minimal paths of length r in \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 minimal paths of length r in \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 $c(\mathcal{C}_s) = \mathcal{B}_s$ and $c(\mathcal{C}_f) = \mathcal{B}_f$, $d_{P_u(\mathbb{C})}(\mathcal{C}_s, \mathcal{C}_f) \geq d_1$. If $d_1 > \Delta_u$, we found a new partition \mathbb{C} with diameter bigger than Δ_u , contradicting our assumption on 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 $c(\mathcal{C}_s) = \mathcal{B}_s$ and $c(\mathcal{C}_f) = \mathcal{B}_f$. For a fixed pair of such elements $\mathcal{C}_s, \mathcal{C}_f$ apply the claim (yielding a sequence of refined partitions corresponding to all minimal paths of length d 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. \square

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 $x \leq y$, and extended by symmetry for $x > y$. 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 for example the following kernel: let $\mathcal{S} = [0, k + 2]$, and let $\varepsilon > 0$ be a sufficiently small constant that ensures that consecutive intervals overlap below (note that by dividing all values by $k + 2$ clearly the same result can be obtained when considering $\mathcal{S} = [0, 1]$). Then define $\kappa(x, y)$ for $x < y$ as follows:

$$\kappa(x, y) = \mathbf{1}_A + \sum_{i=1}^k \mathbf{1}_{B_i \setminus A} \left(x - \frac{(i-1)(k+2)(1-\varepsilon)}{k} \right) \left(\frac{i(k+2)}{k} - y \right),$$

where

$$A = \left\{ (x, y) \in \mathcal{S} \times \mathcal{S} : y - x \leq 1 \right\}$$

and

$$B_i = \left\{ (x, y) \in \mathcal{S} \times \mathcal{S} : \frac{(i-1)(k+2)(1-\varepsilon)}{k} \leq x < y \leq \frac{i(k+2)}{k} \right\}.$$

Clearly, for any partition \mathbb{A} of \mathcal{S} , for both the upper and lower partition graphs of \mathbb{A} the diameter of the corresponding graph is attained for the distance between 0 and $k + 2$. To see that $\Delta_u = k$ consider an arbitrarily fine partition \mathbb{A} : the element of \mathbb{A} containing 0 is connected by an edge in the upper partition graph to the element of \mathbb{A} containing $(k + 2)/k$, this one is then connected by an edge to the element of \mathbb{A} containing $2(k + 2)/k$, and continuing like this, the element containing $(k - 1)(k + 2)/k < k$ is then connected by an edge to the element of \mathbb{A} containing k , and we have found a path of length k . Clearly, in an arbitrarily fine partition these elements of all \mathbb{A} are all different, and no shorter path can be found, since the element of \mathcal{A} containing $i(k + 2)/k$ is not connected by an edge to the element of \mathcal{A} containing only elements bigger than $(i + 1)(k + 2)/k$. By definition of κ is certainly an optimal strategy, and hence $\Delta_u \leq k$. To show that $\Delta_\ell \geq k + 2$, observe that the element of \mathbb{A} containing 0 can possibly be connected by an edge to the element of \mathbb{A} containing 1, but not to an element of \mathbb{A} containing only elements of \mathcal{S} bigger than 1. For ε small enough, $(k + 2)(1 - \varepsilon)/k > 1$, and therefore in the next step the element of \mathbb{A} containing 1 can possibly be connected by an edge at most to the element of \mathbb{A} containing 2, but not to an element containing only elements of \mathbb{A} bigger than 2. Repeating the same argument, for any $i = 2, \dots, k - 1$, the element containing i can possibly be connected by an edge at most to the element of \mathbb{A} containing $i + 1$, but not to an element of \mathbb{A} containing only elements of \mathcal{S} bigger than $i + 1$. By construction of κ this is clearly an optimal strategy, and hence $\Delta_\ell \geq k + 2$. By the first part of the theorem, we then must have $\Delta_u = k$ and $\Delta_\ell = k + 2$. The proof of Theorem 1 is complete.

4. EXPANSION OF NEIGHBORHOODS

Let $G(n, \kappa, p)$ be an inhomogeneous random graph induced by an irreducible kernel κ that is continuous $(\mu \otimes \mu)$ -almost everywhere. In this section we prove that the number of vertices at distance k from a given vertex grows exponentially with k . The main tool we use is concentration inequalities.

Lemma 4. *Let $\mathbb{A} = \{\mathcal{A}_1, \dots, \mathcal{A}_m\}$ be a partition of \mathcal{S} . Let $U \subset V(\mathcal{A}_i) \leq |V(\mathcal{A}_i)|/2$. Let $\mathcal{A}_j \in \mathbb{A}$, with possibly $\mathcal{A}_i = \mathcal{A}_j$ and suppose that $\kappa_\ell(\mathcal{A}_i, \mathcal{A}_j) > 0$. Define*

$$S_{i,j} := n\mu(\mathcal{A}_j) \left(1 - (1 - \kappa_\ell(\mathcal{A}_i, \mathcal{A}_j)p)^{|U|}\right) = \begin{cases} n\mu(\mathcal{A}_j)(1 + o(1)), & p|U| = \omega(1), \\ \Theta(n\mu(\mathcal{A}_j)), & p|U| = \Theta(1), \\ n\mu(\mathcal{A}_j)p|U|(1 + o(1)), & p|U| = o(1). \end{cases}$$

Then

$$\mathbf{P} \left(|N(U) \cap (V(\mathcal{A}_j) \setminus U)| \leq \frac{S_{i,j}}{4} \right) \leq e^{-S_{i,j}/16}.$$

Proof. We write the number of vertices as a sum of indicators

$$|N(U) \cap (V(\mathcal{A}_j) \setminus U)| = \sum_{v \notin U} \mathbf{1}_{X_v \in \mathcal{A}_j} \mathbf{1}_{v \in N(U)}.$$

Note that, for different values of $v \notin U$, the events $[X_v \in \mathcal{A}_j] \cap [v \in N(U)]$ are independent and identically distributed. Fix a vertex $v \notin U$. We want to bound the probability of these events conditionally on $[U \subset V(\mathcal{A}_i)]$. We begin by noting that

$$(2) \quad \mathbf{P}(X_v \in \mathcal{A}_j \mid U \subset V(\mathcal{A}_i)) = \mathbf{P}(X_v \in \mathcal{A}_j) = \mu(\mathcal{A}_j),$$

Moreover, by independence we have

$$\begin{aligned} \mathbf{P}(v \notin N(U) \mid U \subset V(\mathcal{A}_i), X_v \in \mathcal{A}_j) &= \prod_{u \in U} \mathbf{P}(v \notin N(u) \mid X_u \in \mathcal{A}_i, X_v \in \mathcal{A}_j) \\ &\leq (1 - \kappa_\ell(\mathcal{A}_i, \mathcal{A}_j)p)^{|U|}. \end{aligned}$$

Therefore, we can bound

$$(3) \quad \mathbf{P}(v \in N(U) \mid U \subset V(\mathcal{A}_i), X_v \in \mathcal{A}_j) \geq 1 - (1 - \kappa_\ell(\mathcal{A}_i, \mathcal{A}_j)p)^{|U|}$$

Putting equations (2) and (3) together we obtain

$$\begin{aligned} &\mathbf{P}(X_v \in \mathcal{A}_j, v \in N(U) \mid U \subset V(\mathcal{A}_i)) \\ &= \mathbf{P}(X_v \in \mathcal{A}_j \mid U \subset V(\mathcal{A}_i)) \cdot \mathbf{P}(v \in N(U) \mid U \subset V(\mathcal{A}_i), X_v \in \mathcal{A}_j) \\ &\geq \mu(\mathcal{A}_j) \left(1 - (1 - \kappa_\ell(\mathcal{A}_i, \mathcal{A}_j)p)^{|U|}\right) \end{aligned}$$

The proof is complete by applying a Chernoff bound [11] to $|N(U) \cap (V(\mathcal{A}_j) \setminus U)|$ noting that $S_{i,j}/2 \leq \mathbf{E} |N(U) \cap (V(\mathcal{A}_j) \setminus U)|$. \square

We will use the previous lemma iteratively to get concentrations for expansions of the i -th neighborhood. We need a few definitions: let

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

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

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

Let $S := n\|\mathbb{A}\|_\mu\kappa_\ell p/4$ and note that $S_{i,j}/4 \geq S$ for all $\mathcal{A}_j, \mathcal{A}_i \in \mathbb{A}$. Define $t(k) = k$ for $0 \leq k \leq \Phi - 2$ and $t(k) = \Phi - 2$ for $k > \Phi - 2$. Given a partition \mathbb{A} , consider a walk $\mathcal{A}_0, \mathcal{A}_1, \dots, \mathcal{A}_\ell$ in $P_\ell(\mathbb{A})$ (possibly making zigzags and reusing partitions) and let $u \in V(\mathcal{A}_0)$. Define $\Gamma_0(u) = \{u\}$, and for $k \geq 1$, define first recursively the following auxiliary variables $\Gamma'_k(u)$: for $|\Gamma_{k-1}(u)| \leq |V(\mathcal{A}_{k-1})|/2$, let $\Gamma'_{k-1}(u) := \Gamma_{k-1}(u)$,

and otherwise define $\Gamma'_{k-1}(u) \subseteq \Gamma_{k-1}(u)$ to be a randomly chosen subset of $\Gamma_{k-1}(u)$ such that $|\Gamma'_{k-1}(u)| = \lfloor |V(\mathcal{A}_{k-1})|/2 \rfloor$. Define then

$$\Gamma_k(u) = \left\{ v \in V(\mathcal{A}_k) \setminus \Gamma'_{k-1}(u) : N(v) \cap \Gamma'_{k-1}(u) \neq \emptyset \right\}.$$

We can now prove the following lemma.

Lemma 5. *Let ω be a function tending to infinity arbitrarily slowly as $n \rightarrow \infty$, let $p \geq \omega \log n/n$ and let $\ell \leq \omega$. Let \mathbb{A} be a partition of κ and let $\mathcal{A}_0, \mathcal{A}_1, \dots, \mathcal{A}_\ell$ be a walk in $P_\ell(\mathbb{A})$, and let $u \in V(\mathcal{A}_0)$. Then, with probability $1 - o(n^{-\omega})$,*

$$|\Gamma_\ell(u)| \geq S^{t(\ell)}.$$

Proof. Note that $\Gamma_k(u) = N(\Gamma'_{k-1}(u)) \cap (V(\mathcal{A}_k) \setminus \Gamma'_{k-1}(u))$. Define the events $\mathcal{F}_0 = \Omega$ and in general by induction, for $k \geq 1$,

$$\mathcal{F}_k = \bigcap_{i=0}^{k-1} \mathcal{F}_i \cap \left[|\Gamma_k(u)| \geq S^{t(k)} \right].$$

Note that if \mathcal{F}_k holds, we have that $|\Gamma_i(u)| \geq S^{t(i)}$ for all $i = 0, \dots, k$. We will now bound $\mathbf{P}(\mathcal{F}_k^c | \mathcal{F}_{k-1})$ from above and condition under \mathcal{F}_{k-1} . If $|\Gamma'_{k-1}(u)|p = o(1)$, then

$$\begin{aligned} n\mu(\mathcal{A}_k)pS^{t(k-1)}(1/2 + o(1)) &\leq n\mu(\mathcal{A}_k)p|\Gamma'_{k-1}(u)|(1/2 + o(1)) \\ &= n\mu(\mathcal{A}_k)(1 + o(1)) \left(1 - (1 - \kappa_\ell(\mathcal{A}_{k-1}, \mathcal{A}_k)p)^{|\Gamma'_{k-1}(u)|} \right) / 2 \leq \mathbf{E} |\Gamma_k(u)|. \end{aligned}$$

Hence by Lemma 4 we have that

$$\mathbf{P}(\mathcal{F}_k^c | \mathcal{F}_{k-1}) \leq e^{-\Omega(|\Gamma'_{k-1}(u)|)} \leq e^{-\Omega(S^{t(k-1)})}.$$

Otherwise, if $|\Gamma_{k-1}(u)|p = \Omega(1)$, then

$$\Theta(n\mu(\mathcal{A}_k)) = n\mu(\mathcal{A}_k) \left(1 - (1 - \kappa_\ell(\mathcal{A}_{k-1}, \mathcal{A}_k)p)^{|\Gamma_{k-1}(u)|} \right) / 2 \leq \mathbf{E} |\Gamma_k(u)|,$$

and in this case we have

$$\mathbf{P}(\mathcal{F}_k^c | \mathcal{F}_{k-1}) \leq e^{-\Omega(n)}.$$

Putting together these bounds for different values of k we obtain

$$\begin{aligned} \mathbf{P}(\mathcal{F}_\ell^c) &\leq \sum_{k=2}^{\ell} \mathbf{P}(\mathcal{F}_k^c | \mathcal{F}_{k-1}) \\ &\leq \sum_{k=2}^{\Phi-2} e^{-\Omega(S^{t(k-1)})} + \ell e^{-\Omega(n)} \\ &= e^{-\Omega(np)} = n^{-\omega}. \end{aligned}$$

The lemma follows. \square

5. BOUNDING THE DIAMETER

We dedicate this section to the proof of Theorem 2. We continue to write $G(n, \kappa, p)$ for an inhomogeneous random graph induced by an irreducible kernel κ that is continuous $(\mu \otimes \mu)$ -almost everywhere. We break up the proof of the theorem into six 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 .

Lemma 6. *With high probability, $\text{diam } G(n, \kappa, p) \geq \Delta_u$.*

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

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 $u, v \in V(\mathcal{A}_i), V(\mathcal{A}_j)$. Suppose first that in $P_\ell(\mathbb{A})$ there exists $\mathcal{A}_r \in \mathbb{A}$ such that $\mathcal{A}_r, \mathcal{A}_s, \mathcal{A}_j$ is a walk (possibly reusing partitions and making zigzags) of length 2 in $P_\ell(\mathbb{A})$, and such that there exists a walk (possibly reusing partitions and making zigzags) of length $\Phi - 2$ from \mathcal{A}_i to \mathcal{A}_r . Then note that by Lemma 5, with probability at least $1 - o(n^{-\omega})$

$$(4) \quad |\Gamma^{\Phi-2}(u) \cap V(\mathcal{A}_r)| \geq (\|\mathbb{A}\|_\mu \kappa_\ell np/4)^{\Phi-2}.$$

Also, with probability at least $1 - o(n^{-\omega})$, $|\Gamma(v) \cap V(\mathcal{A}_s)| \geq \|\mathbb{A}\|_\mu \kappa_\ell np/4$. Assuming these conditions deterministically, the probability that there is no edge between $\Gamma^{\Phi-1}(u) \cap V(\mathcal{A}_r)$ and $\Gamma(v) \cap V(\mathcal{A}_s)$ is at most

$$(5) \quad (1 - \kappa p)^{\|\mathbb{A}\|_\mu \kappa_\ell np/4)^{\Phi-1} = e^{-\Omega(n^{\Phi-1} p^\Phi)} = n^{-\omega},$$

and hence with probability $1 - n^{-\omega}$ we found a path of length $\Phi < \Delta_\ell$ between u and v . Similarly, if there exists $\mathcal{A}_r \in \mathbb{A}$, where \mathcal{A}_r is such that $\mathcal{A}_r, \mathcal{A}_s, \mathcal{A}_j$ is a walk of length 2, and a walk of length $\Phi - 1$ from \mathcal{A}_i to \mathcal{A}_r , by Lemma 5,

$$|\Gamma^{\Phi-1}(u) \cap V(\mathcal{A}_r)| \geq (\|\mathbb{A}\|_\mu \kappa_\ell np/4)^{\Phi-2},$$

and by the same argument as before, we found a path of length $\Phi + 1 \leq \Delta_\ell$ between u and v . Otherwise, there must exist $\mathcal{A}_r \in \mathbb{A}$, such that $\mathcal{A}_r, \mathcal{A}_s, \mathcal{A}_j$ is a path of length 2, and such that the shortest path between \mathcal{A}_i and any vertex \mathcal{A}_r in $P_\ell(\mathbb{A})$ has length $\Phi \leq \ell \leq \Delta_\ell - 2$. (If that were not the case, then \mathcal{A}_i and \mathcal{A}_j would be at distance bigger than Δ_ℓ , contradicting our assumption on $P_\ell(\mathbb{A})$.) Again by Lemma 5,

$$|\Gamma^\ell(u) \cap V(\mathcal{A}_r)| \geq (\|\mathbb{A}\|_\mu \kappa n(p - p^2)/4)^{\Phi-2},$$

and by the same argument as before, with probability $1 - o(n^{-\omega})$ we get a path of length $\ell + 2 \leq \Delta_\ell$ between u and v . By a union bound over all pairs of vertices, the statement follows. \square

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

Proof. Note that for $p = \Theta(1)$ and $p < 1$, $\Phi = 2$, $\Delta_u \geq 2$ and since $\text{diam } G(n, \kappa, p) \geq \Delta_u$, we have $\text{diam } G(n, \kappa, p) \geq \Phi$ in this case. Otherwise, for $p = o(1)$, using the same X_1, \dots, X_n , 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 $U_{ij} < \kappa(X_i, X_j)p$ or $U_{ij} < p$. Note that $G(n, 1, p)$ is nothing but $G(n, p)$. Recall that by Lemma 1 $\text{diam } G(n, p) = \Phi$. Hence, with high probability we have $\text{diam } G(n, \kappa, p) \geq \text{diam } G(n, p) = \Phi$. \square

Lemma 9. *Suppose $\Phi \geq \Delta_u$. If for every partition \mathbb{A} there exists $i \leq j$, 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 $\text{diam } G(n, \kappa, p) \geq \Phi + 1$.*

Proof. Fix an arbitrary partition \mathbb{A} and let $i \leq j$ as in the statement of the lemma. We will show that with high probability there exist two vertices $u \in V(\mathcal{A}_i)$, $v \in V(\mathcal{A}_j)$ such that v cannot be reached by u in exactly Φ steps. Note first that with probability 1 there is no edge between two vertices $w \in V(\mathcal{A}_k)$, $x \in V(\mathcal{A}_\ell)$ with $\mathcal{A}_k, \mathcal{A}_\ell \notin E(P_u(\mathbb{A}))$, so we may assume this. Note that by Lemma 1, if $(np)^{\Phi-1} - 2n \log n \rightarrow -\infty$, then there exists with high probability a pair of vertices u, v such that $d_G(u, v) \geq \Phi$. Recall that we assume that the stronger condition (1) holds. In this case, observe that with high probability there exist ω disjoint pairs of vertices such that they are all at distance exactly Φ : to show this we use the idea of sprinkling new edges, introduced in [3], combined with the coupling using the same X_1, \dots, X_n described in the previous lemma. Indeed, if there were only constantly many, by adding fresh random edges distributed as $G(n, \kappa, C'p)$ for some large enough constant $C' = C'(p)$, we would on the one hand have a graph that can be coupled to be a subgraph of $G(n, \kappa, (C' + 1)p)$, while still $((C' + 1)pn)^{\Phi-1} - 2n \log n \rightarrow -\infty$. On the other hand, the number of new edges added is at least as big as the ones present before in $G(n, \kappa, p)$, since the number of new edges added is distributed as in $G(n, \kappa, C'p(1 - p))$, and $C'p(1 - p) > p$ for any $p < 1$ and C' sufficiently large. Hence, among the graph corresponding to the new edges there would also be at most constantly many pairs of vertices at distance Φ , and the probability that after adding the new edges, one vertex pair originally at distance Φ would remain at distance Φ , tends to 0, contradicting Lemma 1. Hence, in $G(n, \kappa, p)$ with high probability there are ω disjoint pairs of vertices at distance Φ , and hence, with high probability there will be at least one pair of vertices u, v satisfying $u \in V(\mathcal{A}_i)$ and $v \in V(\mathcal{A}_j)$. Since they cannot be at distance Φ , they have to be at distance at least $\Phi + 1$, yielding the lemma. \square

Lemma 10. *Suppose $\Phi \geq \Delta_\ell$. With high probability $\text{diam } G(n, \kappa, p) \leq \Phi + 1$.*

Proof. As in Lemma 7, consider a partition \mathbb{A} attaining Δ_ℓ , and consider two arbitrary vertices $u, v \in V(\mathcal{A}_i), V(\mathcal{A}_j)$. Observe that in $P_\ell(\mathbb{A})$ there exists $\mathcal{A}_r \in \mathbb{A}$ such that $\mathcal{A}_r, \mathcal{A}_s, \mathcal{A}_j$ is a walk (possibly reusing partitions and making zigzags) of length 2 in $P_\ell(\mathbb{A})$, and such that there exists a walk (possibly reusing partitions and making zigzags) of length $\Phi - 2$ or $\Phi - 1$ from \mathcal{A}_i to \mathcal{A}_r . The proof is then as is in Lemma 7. \square

Lemma 11. *Suppose $\Phi \geq \Delta_\ell$. 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 $\text{diam } G(n, \kappa, p) \leq \Phi$.*

Proof. Consider such a partition \mathbb{A} and consider two vertices $u, v \in V(\mathcal{A}_i), V(\mathcal{A}_j)$: by assumption, in $P_\ell(\mathbb{A})$ there exists $\mathcal{A}_r \in \mathbb{A}$ such that $\mathcal{A}_r, \mathcal{A}_s, \mathcal{A}_j$ is a walk (possibly reusing partitions and making zigzags) of length 2 in $P_\ell(\mathbb{A})$, and such that there exists a walk (possibly reusing partitions and making zigzags) of length exactly $\Phi - 2$ from \mathcal{A}_i to \mathcal{A}_r . The proof is then as is in Lemma 7. \square

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

Proof of Theorem 2. By combining Lemma 6 and 7 (i) follows. The first part of (ii) follows by Lemma 8 and Lemma 7, and the second part follows by adding Lemma 9. Finally, the first part of (iii) follows by Lemma 8 and 10; for the second part Lemma 9 and Lemma 11 is used. \square

Acknowledgements. 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 (CMM, 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. M. Ajtai, J. Komós, and E. Szemerédi, *Largest random component of a k -cube*, *Combinatorica* **2** (1982), no. 1, 1–7.
4. 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].
5. S. Bhamidi, S. Sen, and X. Wang, *Continuum limit of critical inhomogeneous random graphs*, Preprint. arXiv:1404.4118 [math.PR].
6. B. Bollobás, *The diameter of random graphs*, *Transactions of the American Mathematical Society* **267** (1981), no. 1, 41–52.
7. ———, *Random graphs*, Cambridge University Press, 2001.
8. B. Bollobás, S. Janson, and O. Riordan, *The phase transition in inhomogeneous random graphs*, *Random Structures & Algorithms* **31** (2007), 3–122.
9. C. Borgs, J.T. Chayes, L. Lovász, V.T. Sós, and K. Vesztegombi, *Convergent sequences of dense graphs I: Subgraph frequencies, metric properties and testing*, *Advances in Mathematics* **219** (2008), no. 6, 1801–1851.
10. ———, *Convergent sequences of dense graphs II: Multiway cuts and statistical physics*, *Annals of Mathematics* **176** (2012), 151–219.
11. H. Chernoff, *A measure of asymptotic efficiency for tests of a hypothesis based on the sum of observations*, *The Annals of Mathematical Statistics* **23** (1952), 493–507.
12. F. Chung and L. Lu, *The diameter of sparse random graphs*, *Advances in Applied Mathematics* **26** (2001), 257–279.
13. L. Devroye and N. Fraiman, *Connectivity threshold of inhomogeneous random graphs*, *Random Structures & Algorithms* **45** (2014), no. 3, 408–420.

14. V. Klee and D. Larman, *Diameters of random graphs*, Canadian Journal of Mathematics **33** (1981), 618–640.
15. L. Lovász, *Large networks and graph limits*, American Mathematical Society, 2012.
16. L. Lovász and B. Szegedy, *Limits of dense graph sequences*, Journal of Combinatorial Theory, Series B **96** (2006), no. 6, 933–957.
17. A. Nachmias and Y. Peres, *Critical random graphs: Diameter and mixing time*, The Annals of Probability **36** (2008), no. 4, 1267–1286.
18. I. Norros and H. Reittu, *On a conditionally poissonian graph process*, Advances in Applied Probability **38** (2006), no. 1, 59–75.
19. O. Riordan and N. Wormald, *The diameter of sparse random graphs*, Combinatorics, Probability & Computing **19** (2010), no. 5-6, 835–926.
20. B. Söderberg, *General formalism for inhomogeneous random graphs*, Physical review E **66** (2002), no. 6, 066121.

DEPARTMENT OF MATHEMATICS, HARVARD UNIVERSITY, ONE OXFORD STREET, CAMBRIDGE MA 02138, USA

E-mail address: `fraiman@math.harvard.edu`

LABORATOIRE J–A DIEUDONNÉ, UNIVERSITÉ NICE SOPHIA ANTIPOLIS, PARC VALROSE, 06108 NICE CEDEX 02, FRANCE

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