A local central limit theorem for triangles in a random graph

Justin Gilmer* Swastik Kopparty†

December 6, 2014

Abstract

In this paper, we prove a local limit theorem for the distribution of the number of triangles in the Erdos-Renyi random graph $G(n,p)$, where $p \in (0,1)$ is a fixed constant. Our proof is based on bounding the characteristic function $\psi(t)$ of the number of triangles, and uses several different conditioning arguments for handling different ranges of t.

1 Introduction

We will work with the Erdos-Renyi random graph $G(n,p)$. Recall that $G(n,p)$ is the random undirected graph G on n vertices sampled by including each of the $\binom{n}{2}$ possible edges into G independently with probability p. Let S_n be the random variable equal to the number of triangles in $G(n,p)$. Let $\mu_n = \mathbb{E}[S_n] = p^3\binom{n}{3}$ and $\sigma_n = \sqrt{\mathbf{Var}[S_n]} = \Theta(n^2)$ (see the Appendix for an exact calculation of σ_n). Our main result (Theorem 2) states that if p is a fixed constant in $(0,1)$, then the distribution of S_n is *pointwise* approximated by a discrete Gaussian distribution:

$$\Pr[S_n = k] = \frac{1}{\sqrt{2\pi}\sigma_n} e^{-((k-\mu_n)/\sigma_n)^2/2} \pm o(1/n^2). \tag{1}$$

Thus, for every $k \in \mu_n \pm O(n^2)$, we determine the probability that $G(n,p)$ has exactly k triangles, up to a $(1+o(1))$ multiplicative factor.

1.1 Central Limit Theorems

The study of random graphs has over 50 years of history, and understanding the distribution of subgraph counts has long been a central question in the theory. When the edge probability p is a fixed constant in $(0,1)$, there is a classical central limit theorem for the triangle count S_n (as well as for other connected subgraphs). This theorem says that for fixed constants a, b:

$$\left| \Pr\left[a \le (S_n - \mu_n)/\sigma_n \le b\right] - \int_a^b \frac{1}{\sqrt{2\pi}} e^{-t^2/2}\, dt \right| = o(1),$$

*Department of Mathematics, Rutgers University, Piscataway, NJ, USA. jmgilmer@math.rutgers.edu. Research supported by NSF CCF-083727 and CCF-1218711.

†Department of Mathematics & Department of Computer Science, Rutgers University, Piscataway, NJ, USA. swastik.kopparty@rutgers.edu. Research supported in part by a Sloan Fellowship and NSF CCF-1253886.

(in other words, $(S_n - \mu_n)/\sigma_n$ converges in distribution to the standard Gaussian distribution). There are several proofs of the central limit theorem for subgraph counts, as well as some vast generalizations, known today.

The original proofs of the central limit theorem for triangle counts (and general subgraph counts) used the method of moments. This method is based on the fact for all distributions that are uniquely determined by their moments, the convergence of the moments of a sequence of random variables to the moments of the distribution implies convergence in distribution. Application of the moment method to subgraph statistics goes back to Erdos and Renyi's original paper [Erd60]. There were several papers in the 1980's (see [KR83] and [Kar84]) that used the moment method to understand, under increasingly general assumptions, when normalized subgraph counts converge in distribution to the Gaussian distribution. This line of work culminated with a paper by Ruciński [Ruc88] who completely characterized when normalized subgraph counts converge in distribution to the Gaussian distribution.

There are several other approaches to the central limit theorem for triangle counts (and general subgraph counts). Using Stein's method [Ste71], Barbour, Karoński and Ruciński [BKR89] obtained strong quantitative bounds on the error in the central limit theorem for subgraph counts. A technique from the asymptotic theory of statistics, known as U-statistics, was applied by Nowicki and Wierman [NW88] to obtain a central limit theorem for subgraph counts, although, in a slightly less general setting than the theorem of Ruciński. Janson [Jan92] used a similar method with several applications, including central limit theorems for the joint distribution of various graph statistics. None of these techniques, however, seem to be quantitatively strong enough to estimate the point probability mass of the triangle/subgraph counts when the edge probability p is a constant.

1.2 Poisson Convergence

When the edge probability p is small enough (for example, $p \approx c/n$ for triangles), then there are classical results that give good estimates for $\Pr[S_n = k]$. In this regime, the distribution of the subgraph count S_n itself (i.e., without normalization) converges in distribution (and hence pointwise) to a Poisson random variable. Some of the work dedicated to understanding this probability regime goes back to the original paper of Erdos and Renyi [Erd60] who studied the distribution of counts of trees and cycles using the method of moments. Using Chen's [Che75] generalization of Stein's method to the Poisson setting, Barbour [Bar82] showed Poisson convergence for general subgraph counts. In the Poisson setting, the probability mass is concentrated in an interval of constant size and thus all results are "local" in the sense that they bound the point probability mass of these random variables.

For slightly larger $p \in [n^{-1}, O(n^{-(1/2)})]$ (this is the range of p where $\sigma_n = \Theta(\mu_n)$), Röllin and Ross [RR10] showed that the probability mass function for triangle counts (S_n) is close in the ℓ_∞ and total variation metrics to the probability mass function of a translated Poisson distribution (and hence a discrete Gaussian distribution), and asked whether a similar local limit law holds for larger p (See Remark 4.5 of that paper). Our result gives such a law for constant $p \in (0, 1)$ for the ℓ_∞ metric.

1.3 Subgraph counts mod q

Some more recent works studied the distribution of subgraph counts mod q. For example, Loebl, Matousek and Pangrac [LMP04] studied the distribution of S_n mod q in $G(n, 1/2)$. They showed that when $q \in (\omega(1), O(\log^{1/3} n))$, then for every $a \in \mathbb{Z}_q$, the probability that $S_n \equiv a \bmod q$ equals $(1+o(1)) \cdot \frac{1}{q}$. Kolaitis and Kopparty [KK13] also studied this problem in $G(n,p)$ for fixed $p \in (0,1)$. They showed that for every constant q, and every $a \in \mathbb{Z}_q$, the probability that $S_n \equiv a \bmod q$ equals $(1+\exp(-n)) \cdot \frac{1}{q}$. This latter result also generalizes to all connected subgraph counts, and to multidimensional versions for the joint distribution of all connected subgraph counts simultaneously. DeMarco, Kahn and Redlich [DKR14] extended these results of [KK13] to determine the distribution of subgraph counts mod q in $G(n,p)$ for all p. Many of these works use conditioning arguments that are similar to those used here.

1.4 Our result

The above lines of work:

1. the central limit theorem for triangle counts in $G(n,p)$ with p constant,
2. the Poisson local limit theorem for triangle counts in $G(n,p)$ with p close to n^{-1},
3. the uniform distribution of triangle counts mod q in $G(n,p)$ with p constant,

all strongly suggest the truth of our main theorem (Theorem 2): there is a local discrete Gaussian limit law for triangle counts in $G(n,p)$ with p constant.

The high level structure of our proof follows the basic Fourier analytic strategy behind the classical local limit theorem for the sums of i.i.d. integer valued random variables. To show that the distribution of $(S_n - \mu_n)/\sigma_n$ is close pointwise to the discrete Gaussian distribution (as in equation (1)), it suffices to show that their characteristic functions (Fourier transforms) are close in L_1 distance. Specifically, if we define $\psi_n(t) = \mathbb{E}[e^{it(S_n-\mu_n)/\sigma_n}]$, we need to show that:

$$\int_{-\pi\sigma_n}^{\pi\sigma_n} |\psi_n(t) - e^{-t^2/2}| \, dt = o(1).$$

The central limit theorem for triangle counts can be used to bound the above integral in the range $(-A, A)$ for any large constant A. By choosing A large enough, we can bound $\int_{A<|t|<\pi\sigma_n} |e^{-t^2/2}| \, dt$ by an arbitrarily small constant. We are thus reduced to showing that $\int_{A<|t|<\pi\sigma_n} |\psi_n(t)| \, dt = o(1)$. We achieve this using two different arguments. For $A < |t| < n^{0.55}$, we show that $|\psi_n(t)| < \frac{1}{t^{1+\delta}}$ using a conditioning argument, where we first reveal the edges in a set $F \subseteq \binom{[n]}{2}$, and count triangles according to how many edges they have in F. For $n^{0.55} < |t| < \pi\sigma_n$, we show that $|\psi_n(t)|$ is superpolynomially small in t by another conditioning argument, where we partition the vertex set $[n]$ into two sets U and V, first expose all the edges within V, and then consider the increase to the total number of triangles that occurs when we expose the remaining edges.

We conjecture that a similar local discrete Gaussian limit law should hold for the number of copies of any fixed connected graph H in $G(n,p)$, for any p that lies above the threshold probability for appearance of H. It would also be interesting to understand the joint distribution of subgraph counts in $G(n,p)$ for several fixed connected graphs. It seems like there are many interesting questions here and much to be investigated.

2 Notation and Preliminaries

Let $[n]$ denote the set $\{1, 2, \ldots, n\}$. For each positive integer n let K_n be the complete graph on the vertex set $[n]$. The Erdos-Renyi random graph $G(n,p)$ is the graph G with vertex set $[n]$, where for each $e \in \binom{[n]}{2}$, the edge e is present in G independently with probability p. For $e \in \binom{[n]}{2}$, let X_e denote the indicator for the event that edge e is present in G. For $E \subseteq \binom{[n]}{2}$, we will let $\{0,1\}^E$ denote the set of $\{0,1\}$-vectors indexed by E. Likewise $X_E \in \{0,1\}^E$ will be the random vector for which the value on coordinate e is the random variable X_e.

For the rest of the paper $p \in (0,1)$ will be a universal fixed constant. All asymptotic notation will hide constants which may depend on p. We will use S_n to denote the number of triangles in $G(n,p)$ (thus $S_n \in [0, \binom{n}{3}]$). The mean of S_n is $p^3\binom{n}{3}$ and the variance (see Appendix) is $\sigma_n^2 = \Theta(n^4)$. We let R_n denote the normalized triangle count, $R_n \stackrel{\text{def}}{=} \frac{S_n - p^3\binom{n}{3}}{\sigma_n}$.

Fourier inversion formula for lattices: Let Y be a random variable that has support contained in the (shifted) discrete lattice $\mathcal{L} \stackrel{\text{def}}{=} \frac{1}{b}(\mathbb{Z} - a)$ for real numbers a, b. Let $\psi(t) \stackrel{\text{def}}{=} \mathbb{E}[e^{itY}]$ be the characteristic function of Y. Then for all $y \in \mathcal{L}$ it holds that

$$\Pr(Y = y) = \frac{1}{2\pi b} \int_{-\pi b}^{\pi b} e^{-ity}\psi(t)\ dt. \qquad (2)$$

Throughout the paper, for real numbers x we will use $\|x\|$ to denote the distance from x to the nearest integer. We will often apply the following easy bound.

Lemma 1. *Let B be a Bernoulli random variable that is 1 with probability p. Then:*

$$|\mathop{\mathbb{E}}_B[e^{i\theta B}]| \leq 1 - 8p(1-p) \cdot \|\theta/2\pi\|^2 .$$

Proof. Without loss of generality, we may assume that $\theta \in [-\pi, \pi]$. We first state two elementary inequalities:

$$\cos(t) \leq 1 - 8 \cdot \|t/2\pi\|^2 \qquad (\text{for } t \in [-\pi, \pi]) \qquad (3)$$

and

$$\sqrt{1-t} \leq 1 - t/2 \qquad (\text{for } t \leq 1). \qquad (4)$$

Then we have the following,

$$\begin{aligned}
|E[e^{i\theta b}]| &= |p + (1-p)e^{i\theta}| \\
&= \sqrt{p^2 + (1-p)^2 + 2p(1-p)\cos(\theta)} \\
&\leq \sqrt{p^2 + (1-p)^2 + 2p(1-p)\left(1 - 8 \cdot \|\theta/2\pi\|^2\right)} \qquad (\text{applying } (3)) \\
&= \sqrt{1 - 16p(1-p)\|\theta/2\pi\|^2} \\
&\leq 1 - 8p(1-p)\|\theta/2\pi\|^2 \qquad (\text{applying } (4)).
\end{aligned}$$

□

3 Main Result

We now give a formal statement of our main result.

Theorem 2 (Local limit law for triangles in $G(n,p)$)**.** *Let*

$$p_n(x) = \Pr(R_n = x) \quad \text{for } x \in \mathbb{L}_n = \left\{ \frac{k - p^3\binom{n}{3}}{\sigma_n} : k \in \mathbb{Z} \right\}$$

and

$$\mathcal{N}(x) = \frac{1}{\sqrt{2\pi}} e^{-x^2/2} \quad \text{for } x \in (-\infty, \infty).$$

Then as $n \to \infty$,

$$\sup_{x \in \mathbb{L}_n} |\sigma_n p_n(x) - \mathcal{N}(x)| \to 0.$$

Equivalently, we have that for all n, for all $k \in \mathbb{Z}$,

$$\Pr[S_n = k] = \frac{1}{\sigma_n} \cdot \mathcal{N}\left(\frac{k - p^3 \cdot \binom{n}{3}}{\sigma_n} \right) + o\left(\frac{1}{n^2}\right),$$

(where the $o(1)$ term goes to 0 as $n \to \infty$, uniformly in k).

Proof. Let $\psi_n(t) = \mathbb{E}[e^{itR_n}]$. Then the Fourier inversion formula for lattices (equation 2) gives us

$$\sigma_n p_n(x) = \frac{1}{2\pi} \int_{-\pi\sigma_n}^{\pi\sigma_n} e^{-itx} \psi_n(t)\, dt.$$

The standard Fourier inversion formula (for $\mathbb{R}$), along with the well known formula for the Fourier transform of $\mathcal{N}$, gives us:

$$\mathcal{N}(x) = \frac{1}{2\pi} \int_{-\infty}^{\infty} e^{-itx} e^{-t^2/2}\, dt.$$

Therefore,

$$|\sigma_n p_n(x) - \mathcal{N}(x)| \le \int_{-\pi\sigma_n}^{\pi\sigma_n} |\psi_n(t) - e^{-t^2/2}|\, dt + 2\int_{\pi\sigma_n}^{\infty} e^{-t^2/2}\, dt \tag{5}$$

The second term goes to zero as n tends to infinity. Thus, it suffices to show that

$$\int_{-\pi\sigma_n}^{\pi\sigma_n} \left|\psi_n(t) - e^{-t^2/2}\right|\, dt \tag{6}$$

tends to 0.

Let $A > 0$ be a large constant to be determined later. We divide the integral into three regions:

- $R_1 = (-A, A)$

- $R_2 = (-n^{0.55}, -A) \cup (A, n^{0.55})$
- $R_3 = (-\pi\sigma_n, -n^{0.55}) \cup (n^{0.55}, \pi\sigma_n)$

The following three lemmas will help us bound the integral of $\left|\psi_n(t) - e^{-t^2/2}\right|$ in these three regions.

Lemma 3. *Let A be a fixed positive real number. Then*

$$\int_{-A}^{A} \left|\psi_n(t) - e^{-t^2/2}\right| \, dt \to 0$$

as $n \to \infty$.

Lemma 4. *There exists a sufficiently large constant $D = D(p)$ and $\delta > 0$ such that, for all t with $|t| \in (0, n^{0.55}]$,*

$$|\psi_n(t)| \leq D/|t|^{1+\delta}.$$

Lemma 5. *There exists a sufficiently large constant $D = D(p)$ such that, for all t with $|t| \in [n^{0.55}, \pi\sigma_n]$, it holds that*

$$|\psi_n(t)| \leq D/|t|^{50}.$$

We now proceed to bound (6).

By Lemma 3,

$$\int_{R_1} \left|\psi_n(t) - e^{-t^2/2}\right| \, dt \to 0,$$

for any fixed constant A.

For R_2 and R_3 we have the following,

$$\int_{R_2 \cup R_3} \left|\psi_n(t) - e^{-t^2/2}\right| \, dt \leq \int_{R_2 \cup R_3} |\psi_n(t)| \, dt + \int_{R_2 \cup R_3} \left|e^{-t^2/2}\right| \, dt$$

By Lemma 4 and Lemma 5, there exists constants $D = D(p), \delta > 0$ such that, $|\psi_n(t)| \leq \frac{D}{t^{1+\delta}}$ for all n and all t with $|t| \in (0, \pi\sigma_n]$. Therefore,

$$\int_{R_2 \cup R_3} \left|\psi_n(t) - e^{-t^2/2}\right| \, dt \leq \int_{R_2 \cup R_3} \left|\frac{D}{t^{1+\delta}}\right| \, dt + \int_{R_2 \cup R_3} \left|e^{-t^2/2}\right| \, dt.$$

Since $D/|t|^{1+\delta}$ and $e^{-t^2/2}$ both have finite integral over $(-\infty, -1) \cup (1, \infty)$, the last line above can be made smaller than any ϵ for large enough constant $A = A(\epsilon, p)$. □

4 Proof sketch for bounding $|\psi_n(t)|$

In this section we sketch with some more detail the strategy used to bound the characteristic function

$$\psi_n(t) \stackrel{\text{def}}{=} \mathbb{E}[e^{itR_n}].$$

As a warm up, suppose that R_n was the sum of n i.i.d random variables X_i. Then, by independence,

$$\psi_n(t) = \mathbb{E}\left[e^{it\sum_{i=1}^{n} X_i}\right] = \prod_{i=1}^{n} \mathbb{E}\left[e^{itX_i}\right] = \mathbb{E}\left[e^{itX_1}\right]^n.$$

Thus if $\left|\mathbb{E}[e^{itX_1}]\right|$ is bounded sufficiently far from 1, it would follow that $|\psi_n(t)|$ is small. Of course in our case R_n is the sum of dependent random variables, and one does not immediately have the expression decompose as a product. The idea that gets around this issue is to first reveal a subset F of the edges and then, conditioning on the values of the edges in F (assuming some nice event Λ occurs), show that the expectation is small. For certain choices of F the conditional expectation *does* decompose as a product, and thus the estimation becomes easier. If the good event Λ happens with high enough probability, then one has successfully bounded $\psi_n(t)$.

We now show an argument that bounds the $\psi_n(t)$ when $n^{1/2} \ll |t| \ll n$. For starters, suppose F was all the edges of $\binom{[n]}{2}$ except for a perfect matching M, and let X_F denote the indicator vector for the edges in F that appear in G. For $e = \{u, v\} \in M$ let C_e denote the number of paths of length 2 from u to v that appear in G (note any such path must consist of edges in F). Then conditioned on the value of X_F, the expectation becomes

$$\mathbb{E}\left[e^{it(C + \sum_{e \in M} C_e X_e)/\sigma_n}\right]$$

where C denotes the number of triangles that appear consisting only of edges in F. Note that C and the C_e are all constants conditioned on the value of X_F. Also, each $C_e = C_e(X_F)$ is a binomial random variable, and thus each is concentrated around np^2. Thus for a "typical" value of X_F one has (roughly)

$$\begin{aligned}
\left|\mathbb{E}[e^{itR_n} \mid X_F]\right| &= \left|\mathbb{E}[e^{it(\sum_{e \in M} C_e X_e)/\sigma_n}]\right| \\
&\approx \left|\prod_{e \in M} \mathbb{E}[e^{itnp^2 X_e/\sigma_n}]\right| \\
&\leq \left(1 - 8p(1-p)\left\|\frac{tnp^2}{2\pi\sigma_n}\right\|^2\right)^{n/2} && \text{(applying Lemma 1)} \\
&\approx \left(1 - 8p(1-p)\left(\frac{tnp^2}{2\pi\sigma_n}\right)^2\right)^{n/2} && \text{(since } \sigma_n = \Theta(n^2) \text{ and } |t| \ll n) \\
&\approx \left(1 - \Theta(t^2/n^2)\right)^{n/2} \\
&\approx e^{-\Theta(t^2/n)}.
\end{aligned}$$

Thus if $|t| \gg n^{1/2}$ the above will be small.

In Section 6 we push the above analysis to cover the range where $t \leq n^{.55}$. There we instead let M be a bipartite subgraph obtained by taking a disjoint union of k perfect matchings, where k is chosen to depend on t. As above we first reveal all edges in $F \stackrel{\text{def}}{=} \binom{[n]}{2} - M$ and then condition on the value of X_F. We then count triangles according to how many edges are in M, letting C, Y, and

Z denote the number of triangles with 0,1, and 2 edges in M respectively. As before C will be a constant conditioned on X_F, and $Y = \sum_{e \in M} C_e X_e$ is the sum of $nk/2$ independent random variables.

For k large enough, $\mathbb{E}[e^{itY/\sigma_n}]$ will be small conditioned on a "typical" X_F, and the analysis follows just as above because the expectation decomposes as a product. The difficulty now is Z is a degree 2 polynomial in the variables $\{X_e : e \in M\}$, and thus $\mathbb{E}[e^{it(Y+Z)/\sigma_n}]$ does not decompose as a product even after conditioning on X_F. However, by estimating the variance of Z we will find that tZ/σ_n will be tightly concentrated in an interval of size $o(1)$, whereas tY/σ_n will be roughly uniform mod 2π. Therefore, although Z has a complicated dependence on Y, $t(Y+Z)/\sigma_n$ will still be roughly uniform mod 2π and $|\psi_n(t)|$ will be small. It should be noted that we currently do not know how to prove a stronger bound than $|\psi_n(t)| \leq 1/t^{1+\delta}$ in this range (for contrast, the argument with a single perfect matching implies an exponentially small bound). This seems to be a major obstacle for obtaining a stronger quantitative local limit law.

The above argument is not delicate enough to deal with arbitrary t with $|t| \gg n$. In Section 7, we use a different conditioning argument to bound $\psi_n(t)$ for all t with $|t| \in [n^{.55}, \pi\sigma_n]$. This argument is based on partitioning $\binom{[n]}{2}$ into two sets E and F, and then studying the difference between the number of triangles in the two random graphs (X_E, X_F) and (X'_E, X_F) (where X'_E is an independent copy of the random variable X_E).

5 Small $|t|$

In this section we prove Lemma 3.

Lemma 3 (restated). *Let A be a fixed positive real number. Then*

$$\int_{-A}^{A} \left| \psi_n(t) - e^{-t^2/2} \right| \, dt \to 0$$

as $n \to \infty$.

The proof essentially follows from the central limit theorem for triangle counts. We provide some of the details by applying a few standard results from probability theory regarding the method of moments. We begin with some additional preliminaries that we borrow (with minor changes) from Durrett's textbook "Probability: Theory and Examples" [Dur10].

For a random variable X, its *distribution function* is the function $F(x) \stackrel{\text{def}}{=} \Pr[X \leq x]$. A sequence of distribution functions is said to *converge weakly* to a limit F if $F_n(x) \to F(x)$ for all x that are continuity points of F. A sequence of random variables X_n is said to *converge in distribution* to a limit X_∞ (written $X_n \xrightarrow{d} X_\infty$) if their distribution functions converge weakly.

The moment method gives a useful sufficient condition for when a sequence of random variables converge in distribution.

Theorem 6. *Let X_n be a sequence of random variables. Suppose that $\mathbb{E}[X^k]$ has a limit μ_k for each k and*

$$\limsup_{k \to \infty} \mu_{2k}^{1/2k}/2k < \infty;$$

then $X_n \xrightarrow{d} X_\infty$ where X_∞ is the unique distribution with the moments μ_k.

In the Appendix we provide the standard calculation that $\mathbb{E}[R_n^k] \to \mu_k$ for all k, where

$$\mu_k \stackrel{\text{def}}{=} \begin{cases} (k-1)!! & \text{if } k \text{ is even} \\ 0 & \text{if } k \text{ is odd} \end{cases}$$

are the moments of $N(0,1)$. It is easy to check that these moments do not grow too quickly and thus the theorem implies the well known central limit theorem for triangle counts:

$$R_n \xrightarrow{d} N(0,1). \tag{7}$$

Durrett also provides a theorem that relates convergence in distribution to pointwise convergence of characteristic functions.

Theorem 7. Continuity Theorem. *Let $X_n, 1 \le n \le \infty$, be random variables with characteristic functions ϕ_n. If $X_n \xrightarrow{d} X_\infty$, then $\phi_n(t) \to \phi_\infty(t)$ for all t.*

Applying this with (7), we conclude that $\psi_n(t) \to e^{-t^2/2}$ for all t. To finish the proof of Lemma 3 we apply the dominated convergence theorem to conclude, for any fixed A, that

$$\int_{-A}^{A} \left| \psi_n(t) - e^{-t^2/2} \right| \, dt \to 0$$

as $n \to \infty$.

6 Intermediate $|t|$

In this section, we prove Lemma 4.

Lemma 4 (restated). *There exists a sufficiently large constant $D = D(p)$ and $\delta > 0$ such that, for all t with $|t| \in (0, n^{0.55}]$,*

$$|\psi_n(t)| = |\,\mathbb{E}[e^{itR_n}]| = |\,\mathbb{E}[e^{itS_n/\sigma_n}]| \le D/|t|^{1+\delta}.$$

Note that trivially $|\psi_n(t)| \le 1$, and thus the lemma already holds for constant sized t. Thus we will assume that t and n are both bigger than a sufficiently large constant $D(p)$. To make the exposition simpler, we will assume n is even (however, the same argument can be easily seen to apply when n is odd).

We simplify notation by denoting R_n by R, S_n by S, and σ_n by σ. Partition $[n]$ into sets U, V both of size $n/2$ and let $P \subseteq \binom{[n]}{2}$ be the complete bipartite graph between vertex sets U, V. Let $k < \frac{n}{10^{10}}$ be a positive integer to be determined later. Let $M_1, \cdots, M_k \subseteq P$ be pairwise disjoint perfect matchings between U and V. Let $E = M_1 \cup M_2 \cup \cdots \cup M_k$, and let $F = \binom{[n]}{2} \setminus E$.

Recall that for the random graph $G \in G(n,p)$, we use X_e to denote the indicator for whether edge e appears in G. We also use X_E and X_F to denote the $\{0,1\}^E$-valued random variable $(X_e)_{e \in E}$ and the $\{0,1\}^F$-valued random variable $(X_e)_{e \in F}$ respectively. Let $C(X_F), Y(X_E, X_F)$ and $Z(X_E, X_F)$ be random variables that count the number of triangles in $G(n,p)$ which have 0,1, and 2 edges in E respectively (note that, by construction of E, no triangle may have all 3 edges in E). Thus we have $S = C(X_F) + Y(X_E, X_F) + Z(X_E, X_F)$.

We define:

$$\zeta = \mathop{\mathbb{E}}_{X_E, X_F}[Z(X_E, X_F)].$$

We now work towards bounding $\left|\mathbb{E}[e^{itS/\sigma}]\right|$:

$$\left|\mathbb{E}[e^{itS/\sigma}]\right| = \left|\mathop{\mathbb{E}}_{X_E,X_F}[e^{it(C(X_F)+Y(X_E,X_F)+Z(X_E,X_F))/\sigma}]\right|$$
$$= \left|\mathop{\mathbb{E}}_{X_E,X_F}[e^{it(C(X_F)+Y(X_E,X_F)+\zeta)/\sigma} + e^{it(C(X_F)+Y(X_E,X_F)+Z(X_E,X_F))/\sigma} - e^{it(C(X_F)+Y(X_E,X_F)+\zeta)/\sigma}]\right|$$
$$\leq \left|\mathop{\mathbb{E}}_{X_E,X_F}[e^{it(C(X_F)+Y(X_E,X_F)+\zeta)/\sigma}]\right| + \mathop{\mathbb{E}}_{X_E,X_F}\left[\left|e^{it(Z(X_E,X_F))/\sigma} - e^{it\zeta/\sigma}\right|\right]$$

We bound each of the two terms separately in the following two lemmas. We will then use these lemmas to conclude the proof of Lemma 4.

Lemma 8.

$$\left|\mathop{\mathbb{E}}_{X_E,X_F}[e^{it(C(X_F)+Y(X_E,X_F)+\zeta)/\sigma}]\right| \leq e^{-\Theta(t^2k/n)}.$$

Proof. We bound the above expectation by revealing the edges in two stages. We first reveal X_F, and show that with high probability over the choice of X_F, some good event occurs. We then show that whenever this good event occurs, the value of the above expectation over the random choice of X_E is small.

Formally, using the triangle inequality we get:

$$\left|\mathop{\mathbb{E}}_{X_E,X_F}[e^{it(C(X_F)+Y(X_E,X_F)+\zeta)/\sigma}]\right| = \left|\mathop{\mathbb{E}}_{X_F}\left[e^{it(C(X_F)+\zeta)/\sigma} \cdot \mathop{\mathbb{E}}_{X_E}[e^{it(Y(X_E,X_F))/\sigma}]\right]\right| \tag{8}$$
$$\leq \mathop{\mathbb{E}}_{X_F}\left[\left|\mathop{\mathbb{E}}_{X_E}[e^{it(Y(X_E,X_F))/\sigma}]\right|\right]. \tag{9}$$

For $e = \{u, v\} \in E$ and a vector $x_F \in \{0,1\}^F$, we let $Y_e(x_F)$ denote the number of paths of length 2 from u to v consisting entirely of edges $f \in F$ for which $(x_F)_f = 1$ [1]. In this way, for a given x_F, the random variable $Y(X_E, x_F)$ equals $\sum_{e \in E} Y_e(x_F)X_e$.

Define

$$L = \{x_F \in \{0,1\}^F \mid \text{for some } e \in E,\ Y_e(x_F) < np^2/2\}.$$

Let Λ denote the (bad) event that $X_F \in L$.

Claim 9.

$$\Pr_{X_F}[\Lambda] \leq e^{-\Theta(n)}.$$

Proof. Observe that for any given $e \in E$, the distribution of $Y_e(X_F)$ equals $\mathrm{Bin}(m_e, p^2)$, where m_e equals the number of paths of length 2 joining the endpoints of e, and consisting entirely of edges in F. Also note that we have $m_e \geq n - 2k \geq n(1 - 1/10^9)$.

[1] This differs from the exposition in Section 4 (where E is a single perfect matching), in that some length-2 paths between u and v here may contain edges in E. We do not want to count those paths in $Y_e(x_F)$.

By the Chernoff bound, we have:

$$\Pr[\mathrm{Bin}(m_e, p^2) < np^2/2] \leq e^{-np^2(1-p^2)/200}.$$

Taking a union bound over all $e \in E$, we get the claim. $\square$

Next, we show that if we condition on Λ not occurring, then the desired expectation is small.

Claim 10. *For every $x_F \in \{0,1\}^F \setminus L$,*

$$\left| \mathop{\mathbb{E}}_{X_E \in \{0,1\}^E} \left[e^{itY(X_E, x_F)/\sigma} \right] \right| \leq e^{-\Theta(t^2 k/n)}.$$

Proof. Recall that $Y(X_E, x_F) = \sum_{e \in E} Y_e(x_F) X_e$. Thus we have:

$$\begin{aligned}
\left| \mathop{\mathbb{E}}_{X_E} \left[e^{itY(X_E, x_F)/\sigma} \right] \right| &= \left| \mathop{\mathbb{E}}_{X_E} \left[e^{it(\sum_{e \in E} Y_e(x_F) X_e)/\sigma} \right] \right| \\
&= \left| \prod_{e \in E} \mathbb{E} \left[e^{itY_e(x_F) X_e/\sigma} \right] \right| \quad \text{by the mutual independence of } (X_e)_{e \in E} \\
&\leq \prod_{e \in E} \left(1 - 8p(1-p) \left\| \frac{tY_e(x_F)}{2\pi\sigma} \right\|^2 \right) \quad \text{(applying Lemma 1)} \\
&= \prod_{e \in E} \left(1 - 8p(1-p) \cdot \left(\frac{tY_e(x_F)}{2\pi\sigma} \right)^2 \right) \quad (\text{since } t \leq n^{0.55},\ Y_e(x_F) \leq n, \text{ and } \sigma = \Theta(n^2)) \\
&\leq \left(1 - 8p(1-p) \cdot \left(\frac{tnp^2}{4\pi\sigma} \right)^2 \right)^{nk/2} \quad (\text{since } x_F \in L).
\end{aligned}$$

Recall that $\sigma = \sqrt{\frac{n(n-1)(n-2)(n-3)D}{2}}$ for some constant $D \leq 1$. Thus $\frac{tnp^2}{4\pi\sigma} \geq \frac{tp^2}{4\pi n}$. Therefore we may further bound the above expression by:

$$\begin{aligned}
&\leq \left(1 - 8p(1-p) \left(\frac{tp^2}{4\pi n} \right)^2 \right)^{nk/2} \\
&\leq e^{-\frac{t^2 p^5 (1-p) k}{\pi^2 n}} \\
&= e^{-\Theta(t^2 k/n)}
\end{aligned}$$

$\square$

Going back to equation (9) we have

$$\begin{aligned}
\left| \mathop{\mathbb{E}}_{X_E, X_F} [e^{it(C(X_F) + Y(X_E, X_F) + \zeta)/\sigma}] \right| &\leq \mathop{\mathbb{E}}_{X_F} \left[\left| \mathop{\mathbb{E}}_{X_E} [e^{it(Y(X_E, X_F))/\sigma}] \right| \right] \\
&\leq \Pr_{X_F}[X_F \in L] + \max_{x_F \in \{0,1\}^F \setminus L} \left| \mathop{\mathbb{E}}_{X_E} [e^{it(Y(X_E, x_F))/\sigma}] \right| \\
&\leq e^{-\Theta(n)} + e^{-\Theta(t^2 k/n)} \quad \text{(applying claims 9 and 10)} \\
&\leq e^{-\Theta(t^2 k/n)}.
\end{aligned}$$

$\square$

Lemma 11.

$$\mathop{\mathbb{E}}_{X_E,X_F}\left[\left|e^{it(Z(X_E,X_F))/\sigma}-e^{it\zeta/\sigma}\right|\right]\leq O\left(t^{3/2+\delta/2}\left(\frac{k}{n}\right)^{3/2}\right)+O\left(1/t^{1+\delta}\right)$$

Proof. Simplifying the expression we want to bound, we get:

$$\mathop{\mathbb{E}}_{X_E,X_F}\left[\left|e^{it(Z(X_E,X_F))/\sigma}-e^{it\zeta/\sigma}\right|\right]=\mathop{\mathbb{E}}_{X_E,X_F}\left[\left|e^{it(Z(X_E,X_F)-\zeta)/\sigma}-1\right|\right].$$

Thus proving the lemma reduces to proving a concentration bound: namely that $Z(X_E, X_F)$ is close to ζ with high probability. We will bound $\mathbf{Var}_{X_E,X_F}[Z(X_E,X_F)]$ and apply the Chebyshev inequality. This will give the desired concentration.

Let Δ' denote the set of triangles in K_n that have exactly 2 edges in E. For each $r \in \Delta'$, let $T_r(X_E, X_F)$ be the indicator for the triangle r appearing in G. For two triangles $r, s \in \Delta'$, write $r \sim s$ if r and s share an edge. Note for any $r \in \Delta'$ there are at most $6k$ triangles $s \in \Delta'$ for which $r \sim s$.

We have:

$$\begin{aligned}\mathbf{Var}_{X_E,X_F}[Z(X_E,X_F)] &= \sum_{r\in\Delta'}\sum_{s\in\Delta'}\mathbf{Cov}_{X_E,X_F}[T_r(X_E,X_F),T_s(X_E,X_F)]\\ &= \sum_{r\in\Delta'}\sum_{s\sim r}\mathbf{Cov}_{X_E,X_F}[T_r(X_E,X_F),T_s(X_E,X_F)] \quad \text{(using independence)}\\ &\leq |\Delta'|\cdot|6k|\\ &\leq 6nk^3 \qquad \left(\text{since } |\Delta'| = n\tbinom{k}{2}\right)\end{aligned}$$

Applying Chebyshev's inequality with $\lambda = \sqrt{6}\cdot n^{1/2}\cdot t^{1/2+\delta/2}\cdot k^{3/2}$ we have

$$\begin{aligned}\Pr_{X_E,X_F}[|Z(X_E,X_F)-\zeta|>\lambda] &< \frac{\mathbf{Var}_{X_E,X_F}[Z(X_E,X_F)]}{\lambda^2}\\ &< 1/t^{1+\delta}\end{aligned}$$

Recall that $\|x\|$ denotes the distance from real number x to the nearest integer. Let Λ be the (bad) event that $|Z(X_E, X_F) - \zeta| \geq \lambda$. Using the fact that for any real number θ, $|e^{i\theta} - 1| \leq 2\pi \cdot \|\frac{\theta}{2\pi}\|$, we have

$$\begin{aligned}\mathop{\mathbb{E}}_{X_E,X_F}\left[\left|e^{it(Z(X_E,X_F)-\zeta)/\sigma}-1\right|\right] &\leq 2\pi\mathop{\mathbb{E}}_{X_E,X_F}\left[\left\|\frac{t(Z(X_E,X_F)-\zeta)}{2\pi\sigma}\right\|\right]\\ &\leq 2\pi\cdot\Pr[\Lambda^c]\cdot\frac{t\lambda}{2\pi\sigma}+2\pi\cdot\Pr[\Lambda]\cdot\frac{1}{2}\\ &\leq \frac{t\lambda}{\sigma}+\pi\cdot\Pr[\Lambda]\\ &\leq \sqrt{6}\cdot t^{3/2+\delta/2}\cdot\frac{k^{3/2}\cdot n^{1/2}}{\sigma}+\frac{\pi}{t^{1+\delta}}\\ &\leq O\left(t^{3/2+\delta/2}\cdot\left(\frac{k}{n}\right)^{3/2}\right)+O\left(\frac{1}{t^{1+\delta}}\right). \quad (\text{since } \sigma=\Theta(n^2))\end{aligned}$$

This concludes the proof of Lemma 11. □

To conclude the proof of Lemma 4, we apply Lemma 8 and Lemma 11 to get the bound

$$|\mathbb{E}[e^{itS/\sigma}]| \leq e^{-\Theta(t^2k/1000n)} + O\left(t^{3/2+\delta/2} \cdot \left(\frac{k}{n}\right)^{3/2}\right) + O\left(1/t^{1+\delta}\right) \tag{10}$$

It only remains to check that k may be chosen as to make the right hand side of equation (10) bounded by $O(1/t^{1+\delta})$. Set $\delta = 0.01$, and observe that for $\Omega(1) < t < n^{0.55}$, we have the following two relations:

$$\frac{n\log^2(t)}{t^2} = O\left(\frac{n}{t^{5/3+\delta}}\right),$$

$$\frac{n}{t^{5/3+\delta}} = \omega(1).$$

Thus we may choose k to be an integer satisfying:

$$k = \Omega(n\log^2(t)/t^2) \quad \text{and} \quad k = O(n/t^{5/3+\delta}).$$

For such a k we have

$$e^{-\Theta(t^2k/n)} \leq O(1/t^{1+\delta})$$

and

$$t^{3/2+\delta/2}\left(\frac{k}{n}\right)^{3/2} = O(1/t^{1+\delta}).$$

This concludes the proof of Lemma 4.

7 Big $|t|$

In this section we prove Lemma 5.

Lemma 5 (restated). *There exists a sufficiently large constant $D = D(p)$ such that, for all t with $|t| \in [n^{0.55}, \pi\sigma_n]$, it holds that*

$$|\mathbb{E}[e^{itR_n}]| = |\mathbb{E}[e^{itS_n/\sigma_n}]| \leq D/|t|^{50}.$$

The choice of 50 here is arbitrary, in fact the lemma will hold for any fixed constant in place of 50 (as long as $D(p)$ is chosen large enough). We only choose a large number here to remind the reader that the obstacle to a better quantitative local limit law lies in bounding $\psi_n(t)$ for $|t|$ in the range $(0, n^{.55}]$.

As in the previous section, since n is fixed we simplify notation by denoting S_n as S and σ_n as σ.

We will break down the proof into two different cases. Both cases will use a common framework, which we now set up.

Let $[n] = U \cup V$ be a partition of the vertices. Define $X_U = (X_e)_{e\in\binom{U}{2}}$. For every $x_U \in \{0,1\}^{\binom{U}{2}}$, we will show that:

$$\mathbb{E}[e^{itS/\sigma}|X_U = x_U] \leq O\left(\frac{1}{t^{50}}\right).$$

This will imply the desired bound.

From now on, we condition on $X_U = x_U$.

Let $E_U \subseteq \binom{U}{2}$ be the induced graph on U:

$$E_U = \left\{ \{u, u^*\} \in \binom{U}{2} \mid x_{\{u,u^*\}} = 1 \right\}.$$

Note that E_U is determined by x_U and is thus fixed.

For $u \in U$, let $A_u \in \{0,1\}^V$ denote the vector indicating the neighbors of u in V. Thus $A_u = (X_{\{u,v\}})_{v \in V}$.

Let $B \in \{0,1\}^{\binom{V}{2}}$ denote the adjacency vector of $G|_V$. Thus $B = \{X_e\}_{e \in \binom{V}{2}}$.

Note that all the entries of the A_u's and B are independent p-biased Bernoulli random variables. We will now express the number of triangles in G in terms of the A_u's and B (here $\langle \cdot, \cdot \rangle$ denotes the standard inner product over $\mathbb{R}$) :

- Let S_U denote the number of triangles in G with all three vertices in U (note that S_U is determined by x_U and is thus fixed).
- The expression $\sum_{\{u,u^*\} \in E_U} \langle A_u, A_{u^*} \rangle$ counts the number of triangles in G that have exactly two vertices in U.
- Let $P : \{0,1\}^V \to \{0,1\}^{\binom{V}{2}}$ denote the map defined by:

  $$P(r)_{\{u,v\}} = r_u \cdot r_v.$$

  Then $\sum_{u \in U} \langle P(A_u), B \rangle$ counts the number of triangles in G that have exactly two vertices in V.
- Let $Q : \{0,1\}^{\binom{V}{2}} \to \mathbb{N}$ denote the map that sends an adjacency vector b to the number of triangles in the graph represented by b (that is the triangles whose vertices are contained in V).

  Thus $Q(B)$ counts the number of triangles in G with all three vertices in V.

Then we have the following expression for S in terms of the A_u's and B.

$$S = S_U + \sum_{u \in U} \langle P(A_u), B \rangle + \sum_{\{u,u^*\} \in E_U} \langle A_u, A_{u^*} \rangle + Q(B).$$

We now bound $\mathbb{E}[e^{itS/\sigma}]$.

$$
\begin{aligned}
|\mathbb{E}[e^{itS/\sigma}]|^2 &= \left| \mathop{\mathbb{E}}_{(A_u)_{u\in U},B} \left[e^{it(S_U + \sum_{u\in U}\langle P(A_u),B\rangle + \sum_{\{u,u^*\}\in E_U}\langle A_u,A_{u^*}\rangle + Q(B))/\sigma} \right] \right|^2 \\
&\le \mathop{\mathbb{E}}_{B}\left[\left| e^{itQ(B)/\sigma} \cdot \mathop{\mathbb{E}}_{(A_u)_{u\in U}} \left[e^{it(\langle \sum_{u\in U} P(A_u),B\rangle + \sum_{\{u,u^*\}\in E_U}\langle A_u,A_{u^*}\rangle)/\sigma} \right] \right|^2 \right] \\
&\le \mathop{\mathbb{E}}_{B}\left[\left| \mathop{\mathbb{E}}_{(A_u)_{u\in U}} \left[e^{it(\langle \sum_{u} P(A_u),B\rangle + \sum_{\{u,u^*\}\in E_U}\langle A_u,A_{u^*}\rangle)/\sigma} \right] \right|^2 \right] \\
&= \mathop{\mathbb{E}}_{B} \mathop{\mathbb{E}}_{(A_u)_{u\in U}} \mathop{\mathbb{E}}_{(A'_u)_{u\in U}} \left[e^{it(\langle \sum_u P(A_u)-P(A'_u),B\rangle + \sum_{\{u,u^*\}\in E_U}\langle A_u,A_{u^*}\rangle - \sum_{\{u,u^*\}\in E_U}\langle A'_u,A'_{u^*}\rangle)/\sigma} \right] \\
&\qquad \text{(Where for each } u\in U,\ A'_u \text{ is an independent copy of } A_u) \\
&= \mathop{\mathbb{E}}_{(A_u)_{u\in U}} \mathop{\mathbb{E}}_{(A'_u)_{u\in U}} \left[e^{it(\sum_{\{u,u^*\}\in E_U}\langle A_u,A_{u^*}\rangle - \sum_{\{u,u^*\}\in E_U}\langle A'_u,A'_{u^*}\rangle)/\sigma} \cdot \mathop{\mathbb{E}}_{B}\left[e^{it\langle \sum_u P(A_u)-P(A'_u),B\rangle/\sigma} \right] \right] \\
&= \mathop{\mathbb{E}}_{(A_u)_{u\in U}} \mathop{\mathbb{E}}_{(A'_u)_{u\in U}} \left[e^{it(\sum_{\{u,u^*\}\in E_U}\langle A_u,A_{u^*}\rangle - \sum_{\{u,u^*\}\in E_U}\langle A'_u,A'_{u^*}\rangle)/\sigma} \cdot \mathop{\mathbb{E}}_{B}\left[e^{it\langle h_{\mathbf{A},\mathbf{A}'},B\rangle/\sigma} \right] \right].
\end{aligned}
$$

where $\mathbf{A} = (A_u)_{u\in U}$, $\mathbf{A}' = (A'_u)_{u\in U}$, and where $h_{\mathbf{A},\mathbf{A}'} \in \mathbb{Z}^{\binom{V}{2}}$ is given by:

$$
h_{\mathbf{A},\mathbf{A}'} = \sum_{u\in U}(P(A_u) - P(A'_u)).
$$

Observe that for each $e \in \binom{V}{2}$, $(h_{\mathbf{A},\mathbf{A}'})_e$ is distributed as the difference of two binomials of the form $B(|U|, p^2)$ (but the different coordinates of $h_{\mathbf{A},\mathbf{A}'}$ are not independent).

Our goal is to show that with high probability over the choice of $\mathbf{A}, \mathbf{A}'$, we have that:

$$
C \stackrel{\text{def}}{=} \left| \mathop{\mathbb{E}}_{B}\left[e^{it\langle h_{\mathbf{A},\mathbf{A}'},B\rangle/\sigma} \right] \right|
$$

is small in absolute value. This will imply that $\mathbb{E}[e^{itS/\sigma}]$ is small, as desired.

We now achieve this goal for $|t| > n^{0.55}$ using two different arguments (to cover two different ranges of $|t|$), instantiating the above framework with different settings of $|U|$.

7.1 Case 1: $n^{1.001} \le |t| < \pi\sigma$

Suppose $n^{1.001} < |t| < \pi\sigma$. For this argument, we choose $|U| = 1$.

In this case, the coordinates of $h_{\mathbf{A},\mathbf{A}'}$ have the following joint distribution: Let $J \subseteq V$ be a random subset where each $v \in V$ appears independently with probability p. Let J' be an independent copy of J (think of J and J' as two independently chosen neighborhoods of the vertex u). Then the e coordinate of $h_{\mathbf{A},\mathbf{A}'}$ is 1 if $e \subseteq J - J'$, 0 if $e \subseteq J \cap J'$ or $e \subseteq J^c \cap (J')^c$, and -1 if $e \subseteq J' - J$. A Chernoff bound implies that with probability at least $1 - e^{-\Theta(n)}$ the symmetric difference of J and J' will have size at least $np(1-p)/2$. In such a case $h_{\mathbf{A},\mathbf{A}'}$ will have $\binom{np(1-p)/2}{2} = \Theta(n^2)$ non-zero coordinates. From now on we assume that $\mathbf{A}, \mathbf{A}'$ are such that this event occurs (and we call such an $\mathbf{A}, \mathbf{A}'$ "good").

Then we have:

$$
\begin{aligned}
C &= \left| \mathop{\mathbb{E}}_{B} \left[e^{it\langle \sum_u h_{\mathbf{A},\mathbf{A}'}, B\rangle/\sigma} \right] \right| \\
&= \left| \mathop{\mathbb{E}}_{B} \left[\prod_{e \in \binom{V}{2}} e^{it(h_{\mathbf{A},\mathbf{A}'})_e B_e/\sigma} \right] \right| \\
&= \left| \mathop{\mathbb{E}}_{B} \left[\prod_{e \in \binom{V}{2}, (h_{\mathbf{A},\mathbf{A}'})_e \neq 0} e^{it(h_{\mathbf{A},\mathbf{A}'})_e B_e/\sigma} \right] \right| \\
&= \left| \prod_{e \in \binom{V}{2}, (h_{\mathbf{A},\mathbf{A}'})_e \neq 0} \mathop{\mathbb{E}}_{B_e} \left[e^{it(h_{\mathbf{A},\mathbf{A}'})_e B_e/\sigma} \right] \right| \\
&\leq \prod_{e \in \binom{V}{2}, (h_{\mathbf{A},\mathbf{A}'})_e \neq 0} \left(1 - 8p(1-p) \cdot \left\| \frac{t \cdot |(h_{\mathbf{A},\mathbf{A}'})_e|}{2\pi\sigma} \right\|^2 \right) \quad \text{(by Lemma 1)} \\
&\leq \prod_{e \in \binom{V}{2}, (h_{\mathbf{A},\mathbf{A}'})_e \neq 0} \left(1 - 8p(1-p) \cdot \left(\frac{t}{2\pi\sigma} \right)^2 \right) \quad \text{(since } |(h_{\mathbf{A},\mathbf{A}'})_e| \in \{0, \pm 1\} \text{ and } |t| < \pi\sigma) \\
&\leq e^{-\frac{2p(1-p)t^2}{\pi^2\sigma^2} \cdot \Theta(n^2)} \quad \text{since } \mathbf{A}, \mathbf{A}' \text{ is good} \\
&\leq e^{-\Theta(t^2/n^2)} \quad \text{(since } \sigma = \Theta(n^2)).
\end{aligned}
$$

Now we use the fact that $t \geq n^{1.001}$ to conclude that $D \leq \exp(-\Theta(n^{0.002}))$.

Taking into account the probability of $\mathbf{A}, \mathbf{A}'$ being good, we get:

$$
|\mathbb{E}[e^{itS/\sigma}]|^2 < e^{-\Theta(n)} + e^{-\Theta(n^{0.002})} \ll \frac{1}{t^{100}},
$$

as desired.

7.2 Case 2: $n^{0.55} \leq t < n^{1.01}$

Suppose $n^{0.55} < t < n^{1.01}$. For this argument, we choose $|U| = n/2$.

As before, we have:

$$
C = \left| \mathop{\mathbb{E}}_{B} \left[\prod_{e \in \binom{V}{2}} e^{it(h_{\mathbf{A},\mathbf{A}'})_e B_e/\sigma} \right] \right|
$$

Now for each $e \in \binom{V}{2}$, the distribution of $(h_{\mathbf{A},\mathbf{A}'})_e$ is the difference of two binomials of the form $Bin(|U|, p^2)$. Thus, we will typically have $(h_{\mathbf{A},\mathbf{A}'})_e$ around $\sqrt{|U|}$ in magnitude.

For each $e \in \binom{V}{2}$, let Λ_e be the following bad event (depending on $\mathbf{A}, \mathbf{A}'$): $|(h_{\mathbf{A},\mathbf{A}'})_e| \notin (|U|^{0.49}, |U|^{0.51})$. Let $\gamma = \Pr[\Lambda_e]$. By standard concentration and anti-concentration estimates for Binomial distributions, we have that $\gamma \leq 0.1$ (provided n is sufficiently large, depending on p).

Let Λ be the bad event that for more than $|V|^2/4$ choices of $e \in \binom{V}{2}$, the event Λ_e occurs.

Lemma 12. *There is a constant A such that for every k:*

$$\Pr[\Lambda] < \frac{k^{Ak}}{|V|^k}.$$

Proof. Let Z_e be the indicator variable for the event Λ_e. For each e, we have $\mathbb{E}[Z_e] = \gamma \leq 0.1$.

Note that if $e_1, \ldots, e_k$ are pairwise disjoint, then $Z_{e_1}, \ldots, Z_{e_k}$ are mutually independent.

Let $Z = \sum_{e \in \binom{V}{2}} (Z_e - \gamma)$. Note that $\mathbb{E}[Z] = 0$. We will show that $\mathbb{E}[Z^{2k}] \leq k^{O(k)} \cdot |V|^{3k}$. This implies that

$$\Pr[\Lambda] \leq \Pr[Z > |V|^2/8] \leq \Pr[Z^{2k} > (|V|^2/8)^{2k}] \leq \frac{\mathbb{E}[Z^{2k}]}{(|V|^2/8)^{2k}} \leq k^{O(k)} \frac{1}{|V|^k},$$

as desired.

It remains to show the claimed bound on $\mathbb{E}[Z^{2k}]$. We have:

$$\mathbb{E}[Z^{2k}] = \sum_{e_1, \ldots, e_{2k} \in \binom{V}{2}} \mathbb{E}[\prod_{j=1}^{2k} (Z_{e_j} - \gamma)].$$

We call a tuple $(e_1, \ldots, e_{2k}) \in \binom{V}{2}^{2k}$ *intersecting* if for every $i \in [2k]$, there exists $j \neq i$ with $e_j \cap e_i \neq \emptyset$. The key observation is the following: if $(e_1, \ldots, e_{2k})$ is not intersecting, then $\mathbb{E}[\prod_{j=1}^{2k} (Z_{e_j} - \gamma)] = 0$. To see this, suppose $(e_1, \ldots, e_{2k})$ is not intersecting because e_i does not intersect any other e_j. Then we have:

$$\mathbb{E}[\prod_{j=1}^{2k} (Z_{e_j} - \gamma)] = \mathbb{E}[Z_{e_i} - \gamma] \cdot \mathbb{E}[\prod_{j \neq i} (Z_{e_j} - \gamma)] = 0,$$

where the first equality follows from the independence property of the Z_e mentioned above.

Thus, $\mathbb{E}[Z^{2k}] \leq \sum_{(e_1, \ldots, e_{2k}) \text{ intersecting}} 1$. We conclude the proof by counting the number of intersecting tuples $(e_1, \ldots, e_{2k})$. Note that for every intersecting tuple $(e_1, \ldots, e_{2k})$, we have $\left| \bigcup_{j=1}^{2k} e_j \right| \leq 3k$. The number of intersecting tuples where every edge intersects exactly one other edge is $k^{\Theta(k)} n^{3k}$. Notice that every intersecting tuple that is not of this form has $\left| \bigcup_{j=1}^{2k} e_j \right| \leq 3k - 1$. Thus the number of such intersecting tuples is at most $\binom{(3k)^2}{k} \cdot n^{3k-1} = k^{O(k)} \cdot n^{3k-1}$. Thus $\mathbb{E}[Z^{2k}]$ is at most $k^{O(k)} \cdot n^{3k}$, as desired. □

Now suppose Λ does not occur. Then we can bound C as follows:

$$
\begin{aligned}
C &= \left| \mathop{\mathbb{E}}_{B} \left[\prod_{e \in \binom{V}{2}} e^{it(h_{\mathbf{A},\mathbf{A}'})_e B_e/\sigma} \right] \right| \\
&= \prod_{e \in \binom{V}{2}} \left| \mathop{\mathbb{E}}_{B_e} \left[e^{it(h_{\mathbf{A},\mathbf{A}'})_e B_e/\sigma} \right] \right| \\
&\leq \prod_{e \in \binom{V}{2}} \left| \left(1 - 8p(1-p) \cdot \left\| \frac{t(h_{\mathbf{A},\mathbf{A}'})_e}{2\pi\sigma} \right\|^2 \right) \right| \quad \text{(by Lemma 1)} \\
&\leq \prod_{e \in \binom{V}{2} | \neg\Lambda_e} \left| \left(1 - 8p(1-p) \cdot \left\| \frac{t(h_{\mathbf{A},\mathbf{A}'})_e}{2\pi\sigma} \right\|^2 \right) \right| \\
&= \prod_{e \in \binom{V}{2} | \neg\Lambda_e} \left| \left(1 - 8p(1-p) \cdot \left(\frac{t(h_{\mathbf{A},\mathbf{A}'})_e}{2\pi\sigma} \right)^2 \right) \right| \quad \text{(since } t < n^{1.001}, |(h_{\mathbf{A},\mathbf{A}'})_e| < |U|^{0.51}, |U| < n \text{ and } \sigma = \Omega(n^2)) \\
&\leq \prod_{e \in \binom{V}{2} | \neg\Lambda_e} \left| \left(1 - 8p(1-p) \cdot \left(\frac{t|U|^{0.49}}{2\pi\sigma} \right)^2 \right) \right| \quad \text{(since } |(h_{\mathbf{A},\mathbf{A}'})_e| \geq |U|^{0.49}) \\
&\leq e^{-\frac{|V|^2}{8} \cdot 8p(1-p) \cdot \left(\frac{t|U|^{0.49}}{2\pi\sigma} \right)^2}. \quad \text{(since } \Lambda \text{ did not occur)}
\end{aligned}
$$

Now we use the fact that $|U| = |V| = n/2$, that $\sigma = \Theta(n^2)$ and that $n^{0.55} < t$.
Thus $C \leq e^{-\Theta(n^{0.08})}$.
Thus, taking into account the probability of the bad event Λ, we get:

$$
|\mathbb{E}[e^{itS/\sigma}]|^2 \leq O\left(\frac{k^{O(k)}}{n^k} \right) + e^{-\Theta(n^{0.08})} \ll \frac{1}{t^{100}},
$$

(choosing $k = 200$), as desired.

8 Appendix

In this section we compute the moments of the random variable $Z_n \stackrel{\text{def}}{=} S_n - p^3\binom{n}{3}$.

Let Δ denote the set of $\binom{n}{3}$ triangles in K_n. For each $t \in \Delta$ denote X_t to be the indicator of the event that all edges in t appear. We write $t \sim t'$ if triangles t and t' share an edge. Note that if triangles t and t' do not share any edges, the random variables X_t and $X_{t'}$ are independent and

$$
\mathbb{E}[(X_t - p^3)(X_{t'} - p^3)] = 0.
$$

Lemma 13. *Let k be a positive integer. Let $C = C(p)$ be the constant $C(p) \stackrel{def}{=} \mathbb{E}[(X_t - p^3)(X_{t'} - p^3)]$ where t and t' are any two triangles that share exactly one edge. Then if k is odd*

$$
\mathbb{E}[Z_n^k] = O(n^{2k-1})
$$

and if k is even

$$\mathbb{E}[Z_n^k] = \frac{(n)_{2k} C^{k/2} (k-1)!!}{2^{k/2}} + O(n^{2k-1}).$$

Proof. We start with

$$E[Z_n^k] = \sum_{t_1 \in \Delta} \cdots \sum_{t_k \in \Delta} \mathbb{E}\left[\prod_{i=1}^{k} (X_{t_i} - p^3)\right].$$

We say an ordered tuple $(t_1, \cdots, t_k)$ of triangles is *intersecting* if for every i there is a $j \neq i$ for which $t_i \sim t_j$. Note that if $(t_1, \cdots, t_k)$ is not intersecting then there is an i for which the random variable X_{t_i} is independent with X_{t_j} for all $j \neq i$. Furthermore, for such a tuple

$$\mathbb{E}\left[\prod_{i=i}^{k} (X_{t_i} - p^3)\right] = 0.$$

We now split into cases based on the parity of k.

Case k is even:

Given an intersecting tuple we define its *skeleton* to be the subgraph of K_n obtained by taking the union of the triangles t_i. Let H be a graph on $2k$ vertices that consists of $k/2$ connected components, each component being the union of two triangles sharing a single edge (although there are many such graphs H, note they are all isomorphic). We say a tuple $(t_1, \cdots, t_k)$ is *fully paired* if its skeleton is isomorphic to H. We first count the number of fully paired tuples by counting the number of copies of H that appear in K_n times the number of fully paired tuples whose skeleton is H.

To count the copies of H, first note that

$$\binom{n}{4}\binom{n-4}{4} \cdots \binom{n-2k+4}{4} \cdot \frac{1}{(k/2)!} = \frac{(n)_{2k}}{24^{k/2}(k/2)!}$$

counts the number of ways to choose the $k/2$ connected components. Within each component there are 6 choices of the shared edge of the two triangles, after which the two triangles are determined. Thus there are

$$\frac{(n)_{2k} 6^{k/2}}{24^{k/2}(k/2)!} = \frac{(n)_{2k}}{2^k (k/2)!}$$

copies. For each copy there are $k!$ tuples whose skeleton is that copy. Thus the number of fully paired tuples is

$$\frac{(n)_{2k} k!}{2^k (k/2)!} = \frac{(n)_{2k}(k-1)!!}{2^{k/2}}.$$

For a fully paired tuples, the expression $\mathbb{E}\left[\prod_{i=1}^{k} (X_{t_i} - p^3)\right]$ splits as a product of the expectation of each connected component (which are pairwise independent). Thus,

$$\mathbb{E}\left[\prod_{i=1}^{k} (X_{t_i} - p^3)\right] = C^{k/2}. \tag{11}$$

We now quickly argue that the number of intersecting tuples that are not fully paired is $O(n^{2k-1})$. This follows because if a tuple is intersecting but not fully paired, than its skeleton consists of at most $2k-1$ vertices. There are $O(1)$ graphs on a given set of vertices, and given such a graph, there are $O(1)$ tuples whose skeleton is isomorphic to it (k is a constant). Thus there are

$$\sum_{i=3}^{2k-1} O(1)\binom{n}{i} = O(n^{2k-1}) \tag{12}$$

such intersecting graphs.

We then have the following calculation. Let P denote the set of fully paired tuples and Q denote the set of tuples that are intersecting but not fully paired.

$$\begin{aligned}
\mathbb{E}(Z_n^k) &= \sum_{(t_1,\cdots,t_k)} \mathbb{E}\left[\prod_{i=1}^{k}(X_{t_i}-p^3)\right] \\
&= \sum_{(t_1,\cdots,t_k)\in P} \mathbb{E}\left[\prod_{i=1}^{k}(X_{t_i}-p^3)\right] + \sum_{(t_1,\cdots,t_k)\in Q} \mathbb{E}\left[\prod_{i=1}^{k}(X_{t_i}-p^3)\right] \\
&= \frac{(n)_{2k}C^{k/2}(k-1)!!}{2^{k/2}} + O(n^{2k-1}).
\end{aligned}$$

Case k is odd:

Let Q denote the set of intersecting tuples. Note that if k is odd then there are no fully paired tuples of k triangles. Therefore $|Q| = O(n^{2k-1})$ and we have the following:

$$\begin{aligned}
\mathbb{E}(Z_n^k) &= \sum_{(t_1,\cdots,t_k)} \mathbb{E}\left[\prod_{i=1}^{k}(X_{t_i}-p^3)\right] \\
&= \sum_{(t_1,\cdots,t_k)\in Q} \mathbb{E}\left[\prod_{i=1}^{k}(X_{t_i}-p^3)\right] \\
&= O(n^{2k-1}).
\end{aligned}$$

□

Corollary 14. *Let* $\sigma_n^2 \stackrel{def}{=} \textbf{\textit{Var}}[S_n]$ *and let* $R_n \stackrel{def}{=} (S_n - p^3\binom{n}{3})/\sigma_n$. *Then* $\mathbb{E}[R_n^k] \to \mu_k$ *for all* k *fixed, where*

$$\mu_k = \begin{cases} (k-1)!! & \text{if } k \text{ is even} \\ 0 & \text{if } k \text{ is odd} \end{cases}.$$

Acknowledgements

Thanks to Brian Garnett for helpful discussions.

References

[Bar82] AD Barbour. Poisson convergence and random graphs. In *Mathematical Proceedings of the Cambridge Philosophical Society*, volume 92, pages 349–359. Cambridge Univ Press, 1982.

[BKR89] Andrew D Barbour, Michal Karoński, and Andrzej Ruciński. A central limit theorem for decomposable random variables with applications to random graphs. *Journal of Combinatorial Theory, Series B*, 47(2):125–145, 1989.

[Che75] Louis HY Chen. Poisson approximation for dependent trials. *The Annals of Probability*, pages 534–545, 1975.

[DKR14] Bobby DeMarco, Jeff Kahn, and Amanda Redlich. Modular statistics for subgraph counts in sparse random graphs. *arXiv preprint arXiv:1402.2264*, 2014.

[Dur10] Rick Durrett. *Probability: theory and examples*. Cambridge university press, 2010.

[Erd60] Paul Erdos. On the evolution of random graphs. *Publ. Math. Inst. Hungar. Acad. Sci*, 5:17–61, 1960.

[Jan92] Svante Janson. *Orthogonal decompositions and functional limit theorems for random graph statistics*, volume 534. American Mathematical Soc., 1992.

[Kar84] Michal Karonski. *Balanced subgraphs of large random graphs*. UAM, 1984.

[KK13] Phokion G Kolaitis and Swastik Kopparty. Random graphs and the parity quantifier. *Journal of the ACM (JACM)*, 60(5):37, 2013.

[KR83] Michał Karoński and Andrzej Ruciński. On the number of strictly balanced subgraphs of a random graph. In *Graph theory*, pages 79–83. Springer, 1983.

[LMP04] Martin Loebl, Jiří Matoušek, and Ondřej Pangrác. Triangles in random graphs. *Discrete mathematics*, 289(1):181–185, 2004.

[NW88] Krzysztof Nowicki and John C Wierman. Subgraph counts in random graphs using incomplete u-statistics methods. *Annals of Discrete Mathematics*, 38:299–310, 1988.

[RR10] Adrian Röllin and Nathan Ross. Local limit theorems via landau-kolmogorov inequalities. *arXiv preprint arXiv:1011.3100*, 2010.

[Ruc88] Andrzej Ruciński. When are small subgraphs of a random graph normally distributed? *Probability Theory and Related Fields*, 78(1):1–10, 1988.

[Ste71] Charles Stein. Dependent random variables. 1971.