Article:

Benjamini, I., Dikstein, Y., Gross, R. orcid.org/0009-0007-5880-7033 et al. (1 more author) (2025) Randomly twisted hypercubes: between structure and randomness. Random Structures & Algorithms, 66 (1). e21267. ISSN 1042-9832

https://doi.org/10.1002/rsa.21267

Random Structures & Algorithms

WILEY

RESEARCH ARTICLE OPEN ACCESS

Randomly Twisted Hypercubes: Between Structure and Randomness

Itai Benjamini[1] | Yotam Dikstein[2] | Renan Gross[1] | Maksim Zhukovskii[3]

[1]Department of Mathematics, Weizmann Institute of Science, Rehovot, Israel | [2]Department of Computer Science and Applied Mathematics, Weizmann Institute of Science, Rehovot, Israel | [3]School of Mathematical Sciences, Tel Aviv University, Tel Aviv, Israel

Correspondence: Renan Gross (renan.gross@gmail.com)

Received: 22 December 2022 | **Revised:** 25 July 2024 | **Accepted:** 11 October 2024

Funding: I.B. is supported by the Israel Science Foundation. R.G. is supported by the Adams Fellowship Program of the Israel Academy of Sciences and Humanities, the European Research Council and by the Israel Science Foundation. M.Z. is supported in part by the Israel Science Foundation Grant 2110/22.

Keywords: random graphs | semicircle law | shortest path | vertex expansion

ABSTRACT

Twisted hypercubes are generalizations of the Boolean hypercube, obtained by iteratively connecting two instances of a graph by a uniformly random perfect matching. Dudek et al. showed that when the two instances are independent, these graphs have optimal diameter. We study twisted hypercubes in the setting where the instances can have general dependence, and also in the particular case where they are identical. We show that the resultant graph shares properties with random regular graphs, including small diameter, large vertex expansion, a semicircle law for its eigenvalues and no non-trivial automorphisms. However, in contrast to random regular graphs, twisted hypercubes allow for short routing schemes.

1 | Introduction and Construction

The Boolean hypercube Q_n is the graph whose vertex set is $V(Q_n) = \{0,1\}^n$ and whose edge set is $E(Q_n) = \{\{x,y\} \mid x$ and y differ by exactly one coordinate$\}$. One appeasing property of the hypercube graph is its recursive construction: starting with Q_1 as a single edge, Q_n is given by the Cartesian graph product $Q_n = Q_1 \square Q_{n-1}$; essentially, the Cartesian product with an edge amounts to matching together the corresponding vertices of two disjoint copies of Q_{n-1}. See Figure 1 for the first steps of this process.

Generalizing this procedure gives rise to the definition of a *twisted hypercube*, which is obtained by iteratively applying perfect matchings between the vertices of two copies of the original graph.

Definition 1. (σ-twist). Let $G_0 = (V, E_0)$ and $G_1 = (V, E_1)$ be two finite graphs on the same vertex set, and let σ be a permutation of the vertices V. The σ-twist operation, denoted $G_0 \overset{\sigma}{\star} G_1$, produces a graph $G_0 \overset{\sigma}{\star} G_1 = (V', E')$, defined as follows. For $i = 0, 1$, let $V^i = \{(x,i) \mid x \in V\}$ and $F^i = \{\{(x,i),(y,i)\} \mid \{x,y\} \in E_i\}$. Then $G_0 \overset{\sigma}{\star} G_1$ has vertex set

$$V' = V^0 \cup V^1$$

and edge set

$$E' = F^0 \cup F^1 \cup \{\{(x,0),(\sigma(x),1)\} \mid x \in V\}$$

Alternatively, if $A_i \in \mathbb{R}^{m \times m}$ is the adjacency matrix of G_i, and P is the $m \times m$ permutation matrix representing σ, then the adjacency matrix of $G_0 \overset{\sigma}{\star} G_1$ is given by

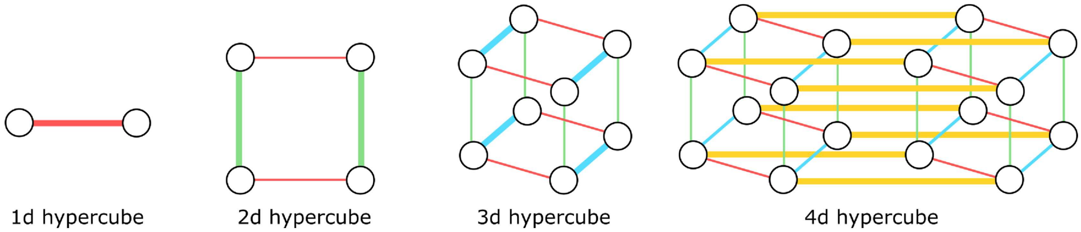

FIGURE 1 | The recursive construction of the hypercube. The color of an edge indicates in which step it was created. Newly created edges are in bold.

$$\begin{pmatrix} A_0 & P \\ P^T & A_1 \end{pmatrix} \tag{1}$$

Definition 2. (Twisted hypercube). Using the σ-twist operation, we define a class of recursively constructed graphs.

1. A *twisted hypercube* graph of generation n, denoted G_n, is defined as follows. For $n = 0$, G_0 is an isolated vertex, labeled by $\emptyset$, and for $n = 1$, G_1 is a single edge, that is, $V(G_1) = \{0, 1\}$ and $E(G_1) = \{\{0, 1\}\}$. For general $n > 1$, let G^0_{n-1} and G^1_{n-1} be two twisted hypercubes of generation $n - 1$, let σ_{n-1} be a permutation on $\{0, 1\}^{n-1}$ vertices, and define
$$G_n = G^0_{n-1} \overset{\sigma_{n-1}}{\star} G^1_{n-1} \tag{2}$$
2. It is possible to consider random permutations in this construction; G_n is then called a *random twisted hypercube*. In this case, the definition requires also specifying the joint distribution of G^0_{n-1} and G^1_{n-1} in the σ-twist operation. For the rest of the article, we assume that all the permutations σ_k are chosen uniformly at random for every k.
3. When the all permutations σ_k, $k = 1, \ldots, n - 1$ are independent of all other permutations, and the two instances G^0_{n-1} and G^1_{n-1} are independent for all n, we call G_n an *independent twisted hypercube*.
4. When the two instances G^0_{n-1} and G^1_{n-1} are identical for all n, we call the graph a *duplicube*. In this case, the graph G_n can be described by a single sequence of permutations $\overline{\sigma} = (\sigma_k)_{k=1}^{\infty}$, where each σ_k is a permutation on $\{0, 1\}^k$: the vertex set is $V(G_n) = \{0, 1\}^n$, and for every $k \in [n]$, the vertex $x = \left(x_1, \ldots x_{k-1}, 0, x_{k+1}, \ldots, x_n\right)$ is connected to $y = \left(\sigma_{k-1}(x_1, \ldots, x_{k-1}), 1, x_{k+1}, \ldots, x_n\right)$. We write $G_n = G_n(\overline{\sigma})$ when we wish to stress the dependence on the permutations. See Figure 2 for the first steps of this process.

The term *twisted cubes* was first introduced in the context of routing in computing networks [1–3]. The idea is that slight modifications to the structure of the hypercube can yield graphs with both better diameter (and so, smaller latency) and better connectivity (and so, better fault-tolerance) than the hypercube. Dudek et al. [4, Definition 2] first introduced randomness to these constructions, and studied independent instances connected by uniform matchings. They named their construction *random twisted hypercubes*. Since our definition generalizes theirs by allowing different joint distributions of matched instances, we have chosen to use the name *random twisted cubes* for the general case, and *independent twisted cubes* for their special case.

Any twisted hypercube G_n is an n-regular graph with $N = 2^n$ vertices. When σ_k is the identity permutation for every k, then G_n is just the Boolean hypercube graph Q_n. The hypercube has diameter n, has poor vertex- and edge-expansion (relative to the fact that its degree grows with the graph size; see Section 2.2), and a random-walk mixing time of order $\Theta(n \log n)$ [5]. Many other geometric and structural properties of the hypercube are known (e.g., distances between vertices [6] and isoperimetric inequalities for various sets [7–9]).

Another well-researched class of n-regular graphs are the uniformly random regular graphs. With probability $1 - o(1)$, a random n-regular graph on 2^n vertices has diameter $\Theta(n/\log n)$, [10] has high edge-expansion [11] and a random-walk mixing time of order $\Theta(n/\log n)$ [12]. Further, its eigenvalues follow a semicircle distribution [13].

For fixed n and $N \to \infty$, the uniform distribution over n-regular graphs on N vertices can be approximated by adding n successive random perfect matchings on N isolated vertices, where the ith matching is uniform over all matchings on previously-unmatched pairs of vertices [14, Theorem 8]. In contrast, consider the random twisted hypercube $\boldsymbol{G}_n$, where all σ_k are uniformly random permutations on $\{0, 1\}^k$, with independence between different k's. It consists of a union of n independent matchings as well, but these matchings are not uniformly random. For example, the last matching is a uniformly random matching only between the two instances of $\boldsymbol{G}_{n-1}$, while earlier matchings consist of a union of smaller matchings; in the case of the duplicube, they consist of *copies* of smaller matchings and therefore have even stronger dependencies between the edges. In this sense, the random twisted hypercube $\boldsymbol{G}_n$ is a hybrid between the structure of the Boolean hypercube and the randomness of a random n-regular graph. It is therefore natural to ask how its various geometric and structural properties compare to those of the hypercube and random n-regular graphs.

Remark 1. An nth iteration duplicube is defined by a single sequence of permutations $\sigma_1, \ldots, \sigma_{n-1}$. To sample such a sequence, one requires approximately $\Theta(n2^n)$ random bits. An independent twisted hypercube, on the other hand, requires $\Theta(n^2 2^n)$ random bits to sample since it is defined using 2^{n-k-1} independent copies of σ_k for every k. As we will see in the next section, despite the fact that it uses less randomness, the structural properties of the duplicube still match those of the independent twisted cube—it has optimal diameter and constant vertex expansion.

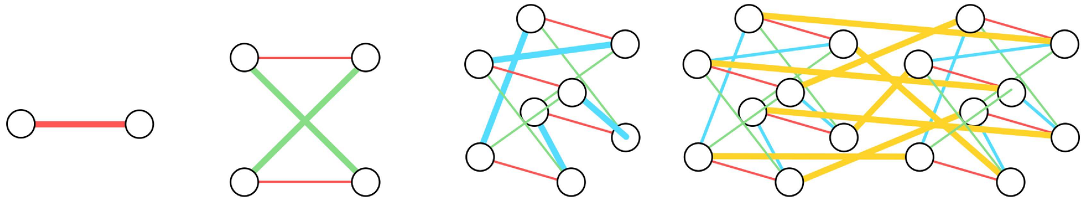

FIGURE 2 | An example of the recursive construction of the duplicube, using random matchings. The color of an edge indicates in which step it was created. Newly created edges are in bold.

2 | Our Results

In this work, we study the diameter, expansion, eigenvalues, and symmetries of a random twisted hypercube $\boldsymbol{G}_n$. All our theorems, except Theorem 1, hold for any twisted hypercube where the matchings σ_k are uniformly random and with independence between permutations of different generations, regardless of the joint distribution of the instances in (2).

2.1 | The Diameter

For a graph $G = (V, E)$, let $d_G : V^2 \to \mathbb{R}$ be the graph distance between two vertices. The diameter of a graph is the maximum distance in the graph, that is, $D(G) := \max\{d_G(x, y) | x, y \in V\}$. An immediate result shows that the diameter of the hypercube Q_n has the worst possible diameter out of all twisted hypercube graphs.

Proposition 1. *For every choice of permutations σ_k, we have $D(G_n) \le D(Q_n) = n$.*

Proof. By induction. For $n = 1$, it is clear. In the general case, let $x, y \in V(G_n)$, and denote $x = (\tilde{x}, x_n)$, $y = (\tilde{y}, y_n)$. If $x_n = y_n$, then $\tilde{x}, \tilde{y} \in V(G_{n-1})$, and $d_{G_n}(x, y) = d_{G_{n-1}}(\tilde{x}, \tilde{y}) \le n - 1$. Otherwise, x is connected to some $(x', 1 - x_n)$, and

$$d_{G_n}(x, y) \le 1 + d_{G_{n-1}}(x', \tilde{y}) \le n$$

□

The following lower bound is also immediate.

Proposition 2. *For every choice of permutations σ_k, we have $D(G_n) \ge (n-1)/\log_2 n$.*

Proof. If G_n has diameter d, then the ball $B(v, d)$ of radius d around any vertex v must contain the entire graph. Since the graph is n-regular, the number of vertices in this ball is smaller than $2n^d$, and we get

$$2^n = |B(v, d)| \le 2n^d$$

yielding

$$d \ge \frac{n-1}{\log_2 n}$$

□

It was shown by Dudek et al. [4] that for the independent twisted hypercube (where the permutations σ_k are chosen uniformly at random, and the instances of $\boldsymbol{G}_{n-1}$ in the σ-twist operation $\boldsymbol{G}_n = \boldsymbol{G}^0_{n-1} \overset{\sigma_{n-1}}{\star} \boldsymbol{G}^1_{n-1}$ are independent), the diameter of $\boldsymbol{G}_n$ is almost surely asymptotic to $n/\log_2 n$. We show that their proof technique carries over to the duplicube as well.

Theorem 1. *Let $\boldsymbol{G}_n$ be the random duplicube. Then $D(\boldsymbol{G}_n) = \frac{n}{\log_2 n} + O(\frac{n}{\log^2 n})$ with probability $\ge 1 - o(2^{-n})$.*

Moreover, we show that regardless of the joint distribution of the two instances of $\boldsymbol{G}_{n-1}$, the diameter is asymptotically better than that of Proposition 1 by at least a $\log\log n / \log\log\log n$ factor. The following theorem is proved in Section 4.2.

Theorem 2. *There exists a constant $C > 0$ such that*

$$D(\boldsymbol{G}_n) \le Cn \frac{\log\log\log n}{\log\log n} \tag{3}$$

with probability $1 - o(2^{-n})$.

Remark 2. The proof of Proposition 1 also gives a simple routing scheme between any two vertices x, y: when at x, let $k \in [n]$ be the largest index such that $x_k \ne y_k$, and go along the edge created by σ_{k-1}. Thus, we always have a local routing scheme which gives a good approximation to the diameter, as well as the average distance between pairs of vertices. Contrast this with general random n-regular graphs, where there is no known local easy way to find an approximation to the minimal path between two vertices.

Remark 3. It might be possible to improve the factor $\frac{\log\log\log n}{\log\log n}$ in Theorem 2 by a more careful analysis of the quantities $\alpha(n)$ and $\beta(n)$ that appear in the theorem's proof. We showed that the diameter of a random twisted hypercube is asymptotically less than that of the Boolean hypercube, yet we have no intuition to the correct diameter.[1]

2.2 | Vertex Expansion

Let $G = (V, E)$ be any graph. For a set $S \subseteq V$, let ∂S be its set of neighbors, that is, $\partial S = \{x \notin S | \exists y \in S \text{ such that } \{x, y\} \in E\}$.

Definition 3. (Vertex expander). Let $0 < \eta < 1$ and $\alpha > 0$, and let $G = (V, E)$ be a graph. A set $S \subseteq V$ is said to have *α-expansion* if $|\partial S| \ge \alpha |S|$. The graph G is an (η, α)-vertex-expander if S has α-expansion for all $S \subseteq V$ of size $|S| \le \eta |V|$.

The hypercube Q_n is not an (η, α)-vertex-expander for any constants $\eta, \alpha > 0$. To see this, fix some constant $\eta > 0$. There is some $\rho > 0$ so that the ball $S = \{x \in \{0,1\}^n | \sum x_i \le \lceil n/2 - \rho\sqrt{n}\rceil - 1\}$ has size $(\eta + o(1))2^n$. However, its boundary is $\partial S = \{x \in \{0,1\}^n | \sum x_i = \lceil n/2 - \rho\sqrt{n}\rceil\}$ and has size at most $\binom{n}{n/2} = 2^n(1+o(1))/\sqrt{\pi n}$. Thus Q_n cannot have an expansion factor α asymptotically larger than $\frac{1}{\sqrt{n}}$ for any constant η. The random twisted hypercube graph, on the other hand, achieves constant expansion with high probability. The following theorem is proved in Section 4.3.

Theorem 3. *For every $\eta \in (0,1)$ there exists a constant $\alpha > 0$ such that*

$$\lim_{n\to\infty} \mathbb{P}[\boldsymbol{G}_n \text{ is a } (\eta,\alpha)\text{-vertex expander}] = 1$$

In fact, the proof of Theorem 3 shows that $\mathbb{P}[\boldsymbol{G}_n \text{is not a } (\eta,\alpha)\text{-vertex expander}] = O(2^{-cn})$ for some constant $c > 0$ that depends on η.

Remark 4. It is also possible to talk about edge expanders, and compare the size of a set S to the number of edges connecting it to ∂S. Both Q_n and G_n are not very good edge expanders (for any choice of permutations σ_k); see Section 3 for more details.

Remark 5. In random d-regular graphs, balls of any constant radius r around an individual vertex are trees with high probability (even when d is logarithmic in the number of vertices). Such sets are very poorly connected-the vertex expansion is of order $1/d^r$ (consider cutting the d-ary tree in half at the central vertex). However, in a random twisted hypercube, a ball of radius r contains G_r as a subgraph, which, by the theorem above, has good vertex expansion with arbitrarily high probability for large r.

2.3 | Eigenvalues

Let $A \in \mathbb{R}^{m\times m}$ be a symmetric matrix, whose eigenvalues are $\lambda_1 \ge \lambda_2 \ge \cdots \ge \lambda_m$. Let $\mu^A := \frac{1}{m}\sum_{i=1}^m \delta_{\lambda_i}$ be the uniform measure over the eigenvalues of A, where δ_s is the Dirac-delta distribution centered at s.

Let $\mathrm{Adj}(Q_n)$ be the adjacency matrix of the hypercube Q_n. The 2^n eigenvalues and eigenvectors of $\mathrm{Adj}(Q_n)$ are well understood; the following is well known [16, section 1.4.6].

Fact 1. For every integer $d \in [0,n]$, the adjacency matrix $\mathrm{Adj}(Q_n)$ has eigenvalue $n-2d$ with multiplicity $\binom{n}{d}$.

In particular, the hypercube's largest eigenvalue is n, while its second largest eigenvalue is $n-2$. Thus, its normalized spectral gap, defined as $\frac{1}{n}(\lambda_1 - \lambda_2)$, is $\frac{2}{n}$. The same gap is achieved for the graphs G_n, regardless of the choice of $\overline{\sigma}$. The following proposition is proved in Section 4.4.

Proposition 3. *Let A_n be the adjacency matrix of G_n. Then $\lambda_1 = n$ and $\lambda_2 = n-2$.*

A consequence of Fact 1 is that $\mu^{\mathrm{Adj}(Q_n)}$ is the probability measure of a $\{\pm 1\}$ Binomial random variable with n trials and success probability $1/2$. By the central limit theorem, we then have that

$$\mu^{\mathrm{Adj}(Q_n)/\sqrt{n}} \to \Gamma$$

weakly, where Γ is the standard Gaussian distribution on $\mathbb{R}$. Unlike the spectral gap, this property is not preserved for the random twisted hypercube graph. In fact, the spectrum of $\mathrm{Adj}(\boldsymbol{G}_n)$ behaves like that of a random n-regular graph.

Theorem 4. *Let $A = \mathrm{Adj}(\boldsymbol{G}_n)$ and $\mu_n = \mu^{A/\sqrt{n}}$. Then the random measure μ_n converges weakly to the semicircle law μ_{circ} in probability, that is, the absolutely continuous measure whose probability density function is*

$$f_{\mathrm{circ}}(x) = \begin{cases} \frac{2}{4\pi^2}\sqrt{4-x^2} & x \in [-2,2] \\ 0 & x \notin [-2,2] \end{cases}$$

The above theorem follows from the following lemma, which states that the number of short cycles in the neighborhood of any vertex in $\boldsymbol{G}_n$ is small. Essentially, this means that $\boldsymbol{G}_n$ is almost locally treelike. For a vertex v and positive integer k, let $\theta(v,k)$ denote the number of cycles of length no more than k containing v, and $B(v,k)$ denote the ball of radius k around v.

Lemma 1. *Let $v \in V_n$ and let $k > 0$ be an integer. There exists a constant $C > 0$ which depends only on k such that the following holds. Let $m_0 > 0$ be an integer, and let*

$$F_v = \bigcup_{u \in B(v,k)} \{\theta(u,k) \le C m_0^{k+1}\}$$

Then

$$\mathbb{P}[F_v] \ge 1 - C2^{-m_0} m_0^{2k+2} n^{2k+1}$$

Theorem 4 and Lemma 1 are proven in Section 4.4.

Remark 6. A classical theorem by McKay [17] states that a regular graph on N vertices has a limiting semicircle law if, for every k, the number of k-cycles in the graph is $o(N)$. This result cannot be directly used in the case of the twisted hypercube: for example, each vertex is guaranteed to be in a 4-cycle, so there are at least $N/4$ 4-cycles in every twisted hypercube (in fact, we conjecture that for the duplicube, for every k, each vertex is in a constant number of k-cycles in expectation). Lemma 1 is the main technical component in our proof of Theorem 4.

2.4 | Asymmetry of G_n

Let $G = (V,E)$ be any graph. A function $\varphi : V \to V$ is called an *automorphism* of G if $\{x,y\} \in E \iff \{\varphi(x), \varphi(y)\} \in E$. The set of all automorphisms of a graph is denoted by $\mathrm{Aut}(G)$, and always contains the trivial automorphism — the identity function Id.

It is well known that for the hypercube, $|\mathrm{Aut}(Q_n)| = n!2^n$, and every automorphism $\varphi(x)$ is of the form $\varphi(x_1,\ldots,x_n) = (x_{\pi(1)} + b_1, \ldots, x_{\pi(n)} + b_n)$ for some permutation $\pi \in S_n$ and

$b \in \{0,1\}^n$. On the other hand, a random regular graph of degree n on 2^n vertices is almost surely asymmetric, that is, almost surely has no non-trivial automorphisms [18, Corollary 3.5]. This is also true for random twisted hypercubes.

Theorem 5.

$$\mathbb{P}[\mathrm{Aut}(\boldsymbol{G}_n) = \{\mathrm{Id}\}] = 1 - O(n^2 2^{-n/20})$$

The proof of Theorem 5 is found in Section 4.5.

2.5 | Different Base Graphs

The twisted hypercube graph is the result of repeatedly applying the σ-twist operation on a single vertex. It is also possible to start with any base graph $G_0 = H$, and define $G_n^H = G_{n-1}^0 \overset{\sigma_{n-1}}{\star} G_{n-1}^1$, where σ_{n-1} is a permutation on $2^{n-1}|V(H)|$ vertices. When each σ_k is a uniformly random permutation on $2^k|V(H)|$ elements, we denote the resulting random graph by $\boldsymbol{G}_n^H$. In this case, we say that $\boldsymbol{G}_n^H$ is a random twisted hypercube with base graph H. None of the main results concerning the diameter, expansion, and eigenvalues are severely affected. This is because as $n \to \infty$, the vast majority of the edges meeting each vertex are those created by the σ-twist operation.

Lemma 2. *Let H be a finite connected graph. Let $\boldsymbol{G}_n^H$ be random twisted hypercube with base graph H. Then there exists a random twisted hypercube $\boldsymbol{G}_n$ and a coupling $(\boldsymbol{G}_n^H, \boldsymbol{G}_n)$ such that*:

1. $D\left(\boldsymbol{G}_n^H\right) \leq D(H)D(\boldsymbol{G}_n)$.
2. If the permutations that define $\boldsymbol{G}_n^H$ are independent then so are the permutations that define $\boldsymbol{G}_n$.
3. If $\boldsymbol{G}_n^H$ is the duplicube with base graph H, then $\boldsymbol{G}_n$ is also a duplicube.

Proof Sketch. Consider the σ-twist operation $\boldsymbol{G}_{k+1}^H = \boldsymbol{G}_k^{H,0} \overset{\sigma_k}{\star} \boldsymbol{G}_k^{H,1}$. By contracting each copy of H in $\boldsymbol{G}_{k+1}^H$ to a single vertex and using Hall's marriage theorem, there exists a set S of 2^k edges induced by σ_k which comprise a perfect matching between the copies of H in the two instances of $\boldsymbol{G}_k^H$. Such a set S naturally induces a permutation on $\{0,1\}^k$. For a given σ_k, let π_k be chosen uniformly at random among all such induced permutations. Then if σ_k is chosen uniformly then π_k is a uniform random permutation on $\{0,1\}^k$. We can use these permutations to generate a graph $\boldsymbol{G}_n$ which is coupled with $\boldsymbol{G}_n^H$ so that $\boldsymbol{G}_n$ is a subgraph of the graph obtained by contracting every copy of H in $\boldsymbol{G}_n^H$ to a single vertex. Thus $D\left(\boldsymbol{G}_n^H\right) \leq D(H)D(\boldsymbol{G}_n)$. □

Thus, both Theorems 1 and 2 continue to hold with only a constant-factor change in the diameter.

Remark 7. If H is not connected, one may simply apply the σ-twist operation several times first until $\boldsymbol{G}_n^H$ is connected (this can be shown to happen with probability tending to 1 as $n \to \infty$), then use that as the base graph.

Corollary 1. (Corollary to Theorem 3). *Let H be a finite graph. For every $\eta \in (0,1)$ there exists a constant $\alpha > 0$ such that the*

$$\lim_{n\to\infty} \mathbb{P}[\boldsymbol{G}_n \text{is a } (\eta,\alpha)\text{-vertex expander}] = 1$$

The proof of the above corollary is essentially identical to that of Theorem 3. The latter only uses the edges created by the last three σ-twist operations, and so the statement still holds for $\boldsymbol{G}_n^H$ as well.

Corollary 2. (Corollary to Theorem 4). *Let H be a finite graph. Let $\mu_n = \mu^{A(\boldsymbol{G}_n^H)/\sqrt{n}}$. Then μ_n converges weakly to the semi-circle law μ_{circ} in probability, that is, the absolutely continuous measure whose probability density function is*

$$f_{\mathrm{circ}}(x) = \begin{cases} \frac{2}{4\pi^2}\sqrt{4-x^2} & x \in [-2,2] \\ 0 & x \notin [-2,2] \end{cases}$$

Proof Sketch. We will assume for simplicity that $|V(H)| = 2^d$ for some integer d. We can couple $\boldsymbol{G}_n^H$ with $\boldsymbol{G}_{n+d}$ by observing that $\boldsymbol{G}_{n+d} = \boldsymbol{G}_n^{\boldsymbol{G}_d}$, and using the same permutations σ_k for $\boldsymbol{G}_n^H$ and $\boldsymbol{G}_n^{\boldsymbol{G}_d}$. Since all the edges due to the permutations are the same for $\boldsymbol{G}_n^H$ and $\boldsymbol{G}_n^{\boldsymbol{G}_d}$, their adjacency matrices differ by no more than $c := |V(H)|$ entries at each row, and all the eigenvalues of the matrix $\Delta = \mathrm{Adj}\left(\boldsymbol{G}_n^H\right) - \mathrm{Adj}\left(\boldsymbol{G}_n^{\boldsymbol{G}_d}\right)$ are bounded by c. Denoting $A := \mathrm{Adj}\left(\boldsymbol{G}_n^{\boldsymbol{G}_d}\right)$, for every integer $k > 0$ we have

$$\begin{aligned}
&\left|\sum_{i=1}^{2^{n+d}} \lambda_i\left(\boldsymbol{G}_n^{\boldsymbol{G}_d}\right)^k - \sum_{i=1}^{2^{n+d}} \lambda_i\left(\boldsymbol{G}_n^H\right)^k\right| \\
&= \left|\mathrm{Tr}\left(A^k\right) - \mathrm{Tr}\left((A+\Delta)^k\right)\right| \\
&= |\mathrm{Tr}(P(A,\Delta))|
\end{aligned}$$

where P is a polynomial of degree k for which in every monomial, A has total degree at most $k-1$. By Von Neumann's trace inequality [19, eq. H.10], if $A_1, \ldots, A_m$ are $N \times N$ symmetric matrices, then

$$\sum_{i=1}^{N} \lambda_i(A_1 \cdots A_m) \leq \sum_{i=1}^{N} \lambda_i(A_1) \cdots \lambda_i(A_m)$$

and so the trace of every monomial in P is bounded above by $c^k \sum_{i=1}^{2^{n+d}} |\lambda_i(A)^{k-1}|$. Thus the difference in the normalized moments of $\boldsymbol{G}_n^{\boldsymbol{G}_d}$ and $\boldsymbol{G}_n^H$ is bounded by

$$\begin{aligned}
&\frac{(n+d)^{-k/2}}{2^{n+d}} \left|\sum_{i=1}^{2^{n+d}} \lambda_i\left(\boldsymbol{G}_n^{\boldsymbol{G}_d}\right)^k - \sum_{i=1}^{2^{n+d}} \lambda_i\left(\boldsymbol{G}_n^H\right)^k\right| \\
&\leq C(k)\frac{(n+d)^{-k/2}}{2^{n+d}} \sum_{i=1}^{2^{n+d}} \left|\lambda_i(A)^{k-1}\right| \\
&\text{(Cauchy–Schwarz)} \\
&\leq C(k)\sqrt{\frac{1}{2^{n+d}} n^{-1} \sum_i \lambda_i(A)^2} \sqrt{\frac{1}{2^{n+d}} n^{-(k-1)} \sum_i \lambda_i(A)^{2k-4}}
\end{aligned}$$

By the proof of Theorem 4, $\frac{1}{2^{n+d}} n^{-(k-2)} \sum_i \lambda_i(A)^{2k-4}$ converges to a constant in probability as $n \to \infty$, which means that the sum on the right-hand side above converges to 0 in probability. This implies that the kth moments of the empirical distribution of the eigenvalues of $\boldsymbol{G}_n^H$ converge to those of the semicircle law. □

Corollary 3. (Corollary to Theorem 5). *Let H be a finite graph. Then*

$$\lim_{n\to\infty} \mathbb{P}\left[\mathrm{Aut}\left(\boldsymbol{G}_n^H\right) = \{\mathrm{Id}\}\right] = 1 - O\left(n^2 2^{-\frac{n}{20}}\right)$$

We omit the proof of Corollary 3 since it is similar to the proof of Theorem 5.

3 | Remarks and Further Directions

1. In [20], Zhu gives a simple-to-define, deterministic sequence of permutations $\overline{\sigma} = (\sigma_k)_{k=1}^{\infty}$ for which the twisted hypercube has asymptotically optimal diameter. What can be said about the expansion, asymmetry, and eigenvalues of this construction? If these properties differ from those of a random twisted hypercube, find a deterministic construction for which they agree.

2. Theorem 1 shows that the random duplicube has diameter $\frac{n}{\log_2 n}(1 + o(1))$. Is it true that for all random twisted cubes the same result holds with high probability?[2]

3. Given a sequence of permutations $\overline{\sigma} = (\sigma_k)_{k=1}^{\infty}$, is there a good local routing scheme for the duplicube $G_n(\overline{\sigma})$ that gives a better approximation than Proposition 1 to the shortest path between two vertices?

4. The twisted-hypercube model can be readily extended to d-dimensional hypergraphs: at every step, create d instances of the current hypergraph, and connect the vertices of the d instances by a perfect matching of d-hyperedges. What can be said about the resultant hypergraph?

5. The graph G_n is, in general, not a good edge-expander. One reason for this are cuts across the matchings σ_k for large k. For example, the two instances of G_{n-1} in G_n each have 2^{n-1} vertices, and are connected by 2^{n-1} edges, giving an isoperimetric ratio of 1. This is not so large for a graph whose degree is n. What can we say about the geometric properties of a set with small edge boundary? For Q_n it is known that sets that have small edge expansion are similar to subcubes [21]. Do non-expanding sets in $\boldsymbol{G}_n$ have similar structure?

6. We show that with high probability $\boldsymbol{G}_n$ is a good vertex expander. However, to our knowledge there is no efficient way to verify that a given graph is a vertex expander: assuming the small-set-expansion hypothesis, it is hard to even approximate the vertex expansion of a graph in polynomial time [22]. Is it possible to exploit the structure of the twisted hypercube to verify this property in time poly(2^n)?

7. By using the same coupon-collector argument as for the hypercube, the mixing time of the lazy simple random walk of any twisted hypercube is $O(n \log n)$. On the other hand, if an nth generation edge is never refreshed, then the random walk stays constrained to one half of the graph, and so the mixing time must also be $\omega(n)$. What is the mixing time for the lazy simple random walk on $\boldsymbol{G}_n$? Is it $o(n \log n)$ with high probability?

8. Is it possible to remove edges from $\boldsymbol{G}_n$ and obtain a (near) constant-degree graph, while maintaining good vertex expansion? Is it possible to approach the vertex expansion of a constant-degree random regular graph in this way?

9. Replace every vertex of Q_n by an n-cycle, obtaining a graph CCC_n; this is known as the cube-connected-cycle [23]. As $n \to \infty$, it is well known that CCC_n converges in the Benjamini–Schramm sense [24] to the lamplighter graph $\mathbb{Z}_2 \wr \mathbb{Z}$. We conjecture that the Benjamini–Schramm limit of the twisted cube-connected-cycle, obtained by replacing every vertex of $\boldsymbol{G}_n$ by an n-cycle, is the 3-regular tree: as $n \to \infty$, a vertex chosen at random from this graph corresponds to a high-generation edge with high probability, and these should not be part of many small cycles.

10. Although Theorem 5 shows that random permutations lead to an asymmetric graph, in general different choices of $\overline{\sigma}$ can lead to different automorphism groups. Can we relate properties of the automorphism group of the duplicube $G_n(\overline{\sigma})$ with properties of $\overline{\sigma}$? In particular, can we find large families of $\overline{\sigma}$ so that $G_n(\overline{\sigma})$ is vertex-transitive? As a non-trivial example, consider the permutations $\sigma_k = \mathrm{Id}$ for $k \neq 2$. There are two essentially different possibilities for σ_2: the first is $\sigma_2 = \mathrm{Id}$, leading to the hypercube Q_n. The second is the matching between a pair of 4-cycles which sends an edge to a non-edge. This leads to a vertex-transitive graph that is not isomorphic to Q_n. Can we find a (perhaps random) vertex-transitive $G_n(\overline{\sigma})$ with improved geometric properties over the hypercube?

11. The argument in Theorem 3 only uses the edges of the last three generations of the twisted hypercube. On the other hand, such an argument could not hold while using only the edges of the last two generations, since the graph induced by the edges of the last two generations is a union of cycles. In fact, we believe that when $\sigma_i = \mathrm{Id}$ for $i < n-2$, the resultant graph does not have constant vertex-expansion with high probability. In light of this, it is natural to ask: for an integer $k > 0$, what are the properties of the twisted hypercube graph, where $\sigma_i = \mathrm{Id}$ for $i < n-k$, and σ_i is uniformly random for $i \geq n-k$? What happens when k grows slowly to infinity with n? This is a natural interpolation between the hypercube Q_n and the completely random twisted hypercube $\boldsymbol{G}_n$.

12. The hypercube Q_n induces a partial order on its vertices in a natural way: $x \leq y$ if $x_i \leq y_i$ for every i. This natural partial order has applications (see e.g., [25, Chapter 6]). The twisted hypercube induces a similar partial order inductively: given the order on G_{n-1}, extend it to G_n by having $(x, 0) < (\sigma_{n-1}(x), 1)$ for all $x \in V_{n-1}$, and by keeping the original order within G_{n-1} in both instances. It can be verified that this is indeed a partial order. What are the properties of this partial order as a function of $\overline{\sigma}$? Are there any

combinatorial applications to the partial order produced by the twisted hypercube?

13. The hypercube Q_n is bipartite, and hence always 2-colorable. On the other hand, the chromatic number χ of random d-regular graphs (of constant degree d) is known to take only one of two possible values with high probability, and satisfies $2\chi \log \chi \approx d$ [26]. What is the chromatic number of $\boldsymbol{G}_n$? It can be shown that it is at least 3 with high probability, and so $\boldsymbol{G}_n$ is in general not bipartite (to see this, consider the case where $\boldsymbol{G}_n$ is bipartite, and look at the probability that σ_n induces an odd cycle in $\boldsymbol{G}_{n+1}$).

4 | Proofs

4.1 | Notation and Definitions

All logarithms are in base e unless otherwise noted. For two sequences $f(n), g(n)$, we write $f = o(g)$ if $\lim |f(n)|/|g(n)| \to 0$. For natural numbers $n, k \in \mathbb{N}$, we set $N = 2^n$ and $K = 2^k$. The set of numbers $1, \ldots, n$ is denoted by $[n]$.

We denote the vertex set of G_n by $V_n := \{0,1\}^n$. For two sets $S_1, S_2 \subseteq V_n$, write $S_1 \sim S_2$ if there are $x \in S_1$ and $y \in S_2$ with $\{x,y\} \in E(G_n)$, and $S_1 \nsim S_2$ otherwise. We say that the edges between two disjoint sets of vertices $A, B \subseteq V_n$ constitute a *matching* if every vertex in $A \cup B$ is adjacent to at most one such edge. The set of neighbors of a vertex $x \in V_n$ are denoted by $N(x)$, and the set of neighbors of a set of vertices $S \subseteq V_n$ by $N(S) = \cup_{x \in S} N(x)$.

Let $x, y \in V_n$. The *generation number* of x and y, denoted by $\gamma(x,y)$, is defined as

$$\gamma(x,y) := n - \max\{1 \leq s \leq n \mid x_i = y_i \forall i \geq s\}$$

that is, n minus the longest common suffix of x and y. If $\{x,y\} \in E(G_n)$ is an edge, then that edge is due to the permutation $\sigma_{\gamma(x,y)-1}$. Supposing that $\gamma(x,y) = k$, we then say that x and y are k-neighbors. Every vertex x has exactly one k-neighbor for every $k \in [n]$; we denote it by $N_k(x)$.

For an integer $r > 0$ and vertex $v \in V_n$, denote by

$$B(v,r) := \{z \in V_n \mid \exists \text{ a path of at most } r \text{ edges from } v \text{ to } z\}$$

the ball of radius r around v, and by

$$B_{<k}(v,r) := \{z \in V_n \mid \exists \text{a path } P \text{ of at most } r \text{ edges from } v \text{ to } z \text{ s.t.} \gamma(x,y) < k \forall \{x,y\} \in E(P)\}$$

the r-neighborhood of v obtained by paths which only use edges of generations smaller than k.

For $1 \leq s < n$, the graph G_n contains multiple disjoint instances of graphs G_s. Indeed, let $z \in \{0,1\}^{n-s}$, and define

$$V_n^z := \left\{ (y,z) \in V_n \mid y \in \{0,1\}^s \right\} \tag{4}$$

Then the induced graph on V_n^z is an instance of G_s (when the construction is deterministic, or in the case of the duplicube, these instances are all isomorphic). The sets V_n^z are disjoint for different z, and partition the vertices of G_n. For a vertex $x \in V_n$, let $I_s(x)$ be the set V_n^z which contains x; it is the set of all vertices in G_n which share a suffix with x of size at least $n - s$, that is,

$$I_s(x) := \{y \in V_n \mid \gamma(x,y) \leq s\}$$

Note that $|I_s(x)| = 2^s$. See Figure 3 for a visual aid. Finally, for a set $S \subseteq V_n$, we denote by $\partial_k S$ the boundary due to the first k generations of edges, that is,

$$\partial_k S = \{x \notin S \mid \exists y \in S, \{x,y\} \in E(G_n), \gamma(x,y) \leq k\}$$

We often write ∂S instead of $\partial_n S$ for brevity.

4.2 | The Diameter

The proof of Theorem 1 resembles the proof of Dudek et al. [4] for the independent twisted hypercube.

Proof of Theorem 1. The main idea of the proof is to show that with high probability, for every $v \in V_n$, the ball around v of radius $\frac{n}{2\log_2 n} + O\left(\frac{n}{\log_2^2 n}\right)$ contains $\geq n2^{n/2}$ vertices in the copy of $\boldsymbol{G}_{n-1}$ which contains v. If this holds, then the diameter is $\frac{n}{\log_2 n} + O\left(\frac{n}{\log_2^2 n}\right)$: for every $v \in V_n^0, u \in V_n^1$, denote by S_v^{n-1}, S_u^{n-1} the balls around v, u in $\boldsymbol{G}_{n-1}$. The probability that S_v^{n-1}, S_u^{n-1} are connected by the last permutation is

$$1 - \frac{\binom{2^{n-1} - n2^{n/2}}{n2^{n/2}}}{\binom{2^{n-1}}{n2^{n/2}}} \geq 1 - \left(1 - \frac{n2^{(n-1)/2}}{2^{n-1}}\right)^{n2^{(n-1)/2}} \geq 1 - e^{-n^2}$$

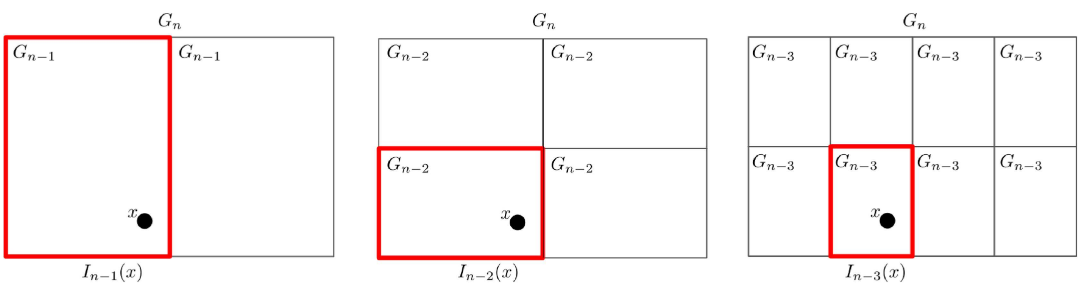

FIGURE 3 | Each large rectangle represents the same graph G_n, with the same vertex x highlighted. The partition into instances of G_s is shown for $s \in \{n-1, n-2, n-3\}$ while highlighting $I_s(x)$.

By a union bound, with high probability every two such balls are connected, so we can find a path of length $\frac{n}{\log_2 n} + O\left(\frac{n}{\log_2^2 n}\right)$ from v to u. If v, u are in the same V_n^i then we use the path from v to $\sigma_{n-1}(u)$ and go from $\sigma_{n-1}(u)$ to u via one additional edge.

Hence we shall show that for any fixed $v \in V$, the ball of radius $\frac{n}{2\log_2 n} + O\left(\frac{n}{\log_2^2 n}\right)$ contains $\geq n2^{n/2}$ vertices in the copy of $\boldsymbol{G}_{n-1}$ which contains v, with probability $1 - o\left(n2^{-2n}\right)$. Beginning with an empty graph on $V = \{0,1\}^n$, we add edges by revealing the values of the permutations σ_k one-by-one, as follows. Set $q = 0.9n$.

1. Initiate a queue Q and insert $v \in Q$.
2. While $Q \neq \emptyset$:
   Take out the first $u \in Q$ and for every $i = q, q+1, \dots, n-2$:
   a. Let $u = (u_1, b, u_2)$ so that $u_1 \in \{0,1\}^i$. If $b = 0$, set $\pi = \sigma_i$, and otherwise set $\pi = \sigma_i^{-1}$; then, if $\pi(u_1)$ wasn't previously revealed, reveal it. This is called the *revealing step*.
   b. For every $u' \in \{0,1\}^{n-1-i}$, we add all edges $\{(u_1, b, u'), (\pi(u_1), 1-b, u')\}$.
   c. For every vertex w that was connected to u and was not previously added to Q, we add $w \in Q$.
3. After Q is empty reveal all other edges in an arbitrary order.

We note that when revealing an entry $\sigma_i(u)$ or $\sigma_i^{-1}(u)$, we in fact add 2^{n-1-i} edges to $\boldsymbol{G}_n$ that come from the different copies of the ith generation duplicube. We say that a vertex u is *discovered* at step k, if one of the edges revealed in the kth step is the first edge that is adjacent to u (where k refers to the number of times we have done step (2a)). Let G' be the subgraph of $\boldsymbol{G}_n$ whose edges are only the edges of generations $i \geq q$. Let S_j be the set of vertices $u \in V$ so that $d_{G'}(v, u) = j$ and so that if u was discovered at step k, then no vertex with the same $0.9n$-prefix of u was discovered previously. Finally, fix r_0 to be the smallest integer so that $n^{r_0} \geq 2^{0.1n}$ and let r_1 be the smallest integer so that $(n/1000)^{r_1} \geq 1000n2^{n/2}$. Clearly $r_0 = \frac{0.1n}{\log_2 n} + O(1)$ and $r_1 = \frac{n}{2\log_2 n} + O\left(\frac{n}{\log_2^2 n}\right)$.

We will analyze the growth of $F(j) = \left|S_j\right|$ separately for $j \leq r_0$ and $r_0 < j \leq r_1$ starting with $j \leq r_0$. We say that the kth step is *bad* if $u = (u_1, b, u_2)$ was the vertex taken out of Q, the value $x = \sigma_i^{\pm 1}(u_1)$ was revealed, and there exists a vertex w whose prefix is x that was discovered in a previous step. At the first phase, we will show that there are very few bad steps. First we calculate a bound on the number of steps m while the distance between v and the vertex that was taken out of the queue in the step is of distance $j \leq r_0$. Namely,

$$m \leq (n-q)\sum_{j=1}^{r_0}(n-q-1)^j \leq 2n(0.1n)^{r_0} \leq 2^{0.11n}$$

Moreover, for the ℓth step for $\ell \leq m$, the probability that ℓ is a bad step is at most

$$p_\ell \leq \frac{\ell(n-q)+1}{2^{0.9n} - 1 - \ell(n-q)}$$

where the numerator is the number of previously discovered vertices (which upper bounds the number of prefixes discovered), and the denominator is the number of choices left. We note that the probability of choosing a given prefix is not uniform, but if a given prefix has already been chosen it only decreases its probability to be chosen again. This is at most

$$p_\ell \leq \frac{n2^{0.11n}}{2^{0.899n}} \leq 2^{-0.78n}$$

As this bound is uniform for all ℓ and the same bound holds true for the conditional probability subject to any way of revealing the first $(\ell - 1)$ edges, the probability p of having c bad steps in the first phase is at most

$$p \leq \binom{m}{c} 2^{-0.78nc} \leq 2^{0.11nc - 0.78nc} = 2^{-0.67cn}$$

Taking $c = 4$ we get that this probability is $o\left(2^{-2n}\right)$. Let k be a step where the vertex taken out of Q is in S_{j-1}. If this kth step is not bad, then S_j grows by 1 due to a new vertex discovered, and if there is a bad edge then it reduces the size by at most 2 (since there were at most two prefixes involved in choosing the bad step). Therefore,

$$\begin{aligned}
&F(1) \geq n - q - 2c \geq (n/1000)\\
&F(2) \geq (F(1) - 2c)(n-q) \geq 0.1n(0.001n - 8) \geq (n/1000)^2\\
&\cdots\\
&F(r_0) \geq (F(r_0 - 1) - 2c)(n-q)\\
&\quad \geq \left((n/1000)^{r_0 - 1} - 8\right)0.1n \geq (n/1000)^{r_0}
\end{aligned}$$

During the second phase, we don't expect there to be no bad steps, but as the set S_j is already quite large, we expect that S_j will still grow by an $\frac{n}{1000}$-factor. Indeed, conditioned on $F(j) \geq (n/1000)^j \geq 2^{0.005n}$, we show that $F(j+1) \geq \frac{n}{1000}F(j)$ with probability $\geq 1 - o\left(2^{-2n}\right)$. When all these events occur, we can conclude that $F(r_1) \geq n2^{n/2}$ with probability $1 - o\left(n2^{-2n}\right)$. Fix $j > r_0$ and let X be a random variable counting the number of bad steps exposed from the vertices of S_j. The number of new vertices we discovered up to this step is at most n^j (and this is also a bound for the number of prefixes discovered). For every step in this phase, the probability that it is bad is at most

$$p' \leq \frac{n^j}{2^{0.9n} - n^j}$$

and as $n^j \leq n^{r_1} \leq 2^{0.51n}$, this is at most $\frac{n^j}{2^{0.6n}}$. We can bound X from above with a $\left(F(j), \frac{n^j}{2^{0.6n}}\right)$-binomially distributed random variable. Thus $\mathbb{E}[X] \leq F(j)\frac{n^j}{2^{0.6n}} = o(F(j))$. Furthermore, by Chernoff's bound on binomial variables

$$\mathbb{P}[X \geq F(j)/10000] \leq e^{-\Omega(F(j))} \leq o\left(2^{-2n}\right)$$

When this event doesn't occur, then

$$\begin{aligned}
F(j+1) \geq (F(j) - 2X)(n-q) &\geq 0.9998\\
\cdot F(j) \cdot 0.1n &\geq F(j)\frac{n}{1000}
\end{aligned}$$

as required.

Remark. As mentioned before, the proof of Theorem 1 is an adaptation of the proof in Dudek et al. [4] for the small diameter of the independent twisted hypercube. The main change is that in this proof, revealing the value of a permutation $\sigma_k(u_1)$ (for $u_1 \in \{0,1\}^k$) reveals many edges between $(u_1, 0, u_2)$ to $(\sigma_k(u_1), 1, u_2)$. Hence, instead of accounting for the number of vertices already discovered, we account for the number of prefixes already discovered.

Proof of Theorem 2. The proof is by induction. For the base cases, let $\gamma > 0$ to be chosen later, and let $k^* = \gamma \log n \log\log n$. by Proposition 1, for all $n \le k^*$ we have

$$\begin{aligned} D(\boldsymbol{G}_n) &\le n \le \log\log k^* \frac{n}{\log\log n} \\ &\le C \log\log\log n \frac{n}{\log\log n} \end{aligned}$$

and so (3) holds with probability 1 for C large enough (depending on γ).

For the induction step, let $n > k^*$. By increasing γ, we may assume that n is larger than any given global constant; this will ensure that inequalities which hold only when n is large enough indeed hold. For an integer $k \ge 1$ and $z \in \{0,1\}^n$, let $\boldsymbol{G}_k^z$ be the induced graph on $I_k(z)$; this is an instance of $\boldsymbol{G}_k$. Denote by E_k^z the event that $D\left(\boldsymbol{G}_k^z\right) \le Ck \frac{\log\log\log k}{\log\log k}$, and assume that E_k^z holds for every $k = 1, \dots, n-1$ and every $z \in \{0,1\}^n$. Let $x, y \in V_n$. If $x_n = y_n$, that is, the two vertices are in the same half of the graph $\boldsymbol{G}_n$, then by the induction hypothesis, $D\left(\boldsymbol{G}_{n-1}^x\right) \le C(n-1)\frac{\log\log\log(n-1)}{\log\log(n-1)}$, and we certainly have $d_{\boldsymbol{G}_n}(x,y) \le Cn\frac{\log\log\log n}{\log\log n}$.

For the case $x_n = 1 - y_n$, that is, the two vertices are in opposite sides of the graph $\boldsymbol{G}_n$, we'll show that for a not-too-large radius, the spheres around x and y contain enough vertices, so that with high probability there is an edge between them induced by σ_{n-1}.

Given a vertex $v \in V_n$ and any integers $t \ge s \ge 0$, let $M(v,s,t) = \{(N_k(v), k) | s \le k \le t\}$ be the set of neighbors of v whose edge to v was added at times $s \le k \le t$, along with their generation number. Note that for $(z,k) \in M(v,s,t)$, the set $I_k(z)$ is contained in $I_{t+1}(v)$, and that since each k-neighbor is added at a different generation, the sets $\{I_k(z)\}_{(z,k)\in M(v,s,t)}$ are all mutually disjoint (see Figure 4). We can therefore iteratively apply the function $M(v,s,t)$ to obtain a large set of disjoint vertices.

More formally, let $s, \ell > 0$ be integers, and consider a subset $S_\ell(x)$ of the sphere of radius ℓ around x, defined as follows:

$$\begin{aligned} S_0(x) &= \{(x, n-1)\} \\ S_i(x) &= \bigcup_{(z,t)\in S_{i-1}(x)} M(z, s, t-1) \end{aligned}$$

By the remark above, the sets $\{I_s(z)\}_{(z,t)\in S_\ell(x)}$ are all disjoint, and so the set $U(x) := \cup_{(z,t)\in S_\ell(x)} I_s(z)$ has cardinality $2^s|S_\ell(x)|$. Define $S_\ell(y)$ and $U(y)$ similarly. Write the values of s and ℓ as $s = \frac{n}{2} - \frac{1}{2}\alpha(n)$ and $\ell = \frac{n}{\beta(n)}$, for some functions $\alpha, \beta : \mathbb{N} \to \mathbb{N}$ to be chosen later. Assuming that there is a vertex $u \in U(x)$ which is connected to $v \in U(y)$, the distance between x and y can be bounded as follows:

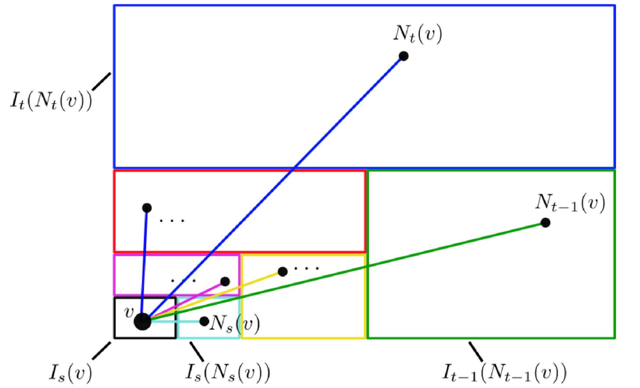

FIGURE 4 | The entire rectangle represents the graph $I_{t+1}(v)$. Each neighbor $N_k(v)$ is contained in $I_k(N_k(v))$, and these $I_k(N_k(v))$ are all disjoint.

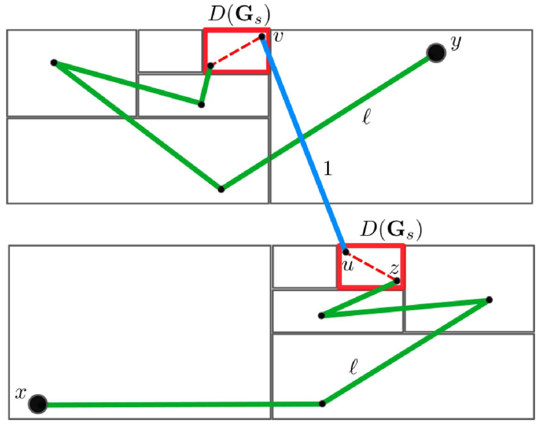

FIGURE 5 | If $U(x)$ is connected to $U(y)$, we have a path from x to y. The red dotted lines represent an optimal path within $\boldsymbol{G}_s$.

$$d_{\boldsymbol{G}_n}(x,y) \le 2\ell + D\left(\boldsymbol{G}_s^u\right) + D\left(\boldsymbol{G}_s^v\right) + 1 \tag{5}$$

where ℓ bounds the distance to go from x to a vertex z in $S_\ell(x)$, $D\left(\boldsymbol{G}_s^u\right)$ bounds the distance from z to u, and 1 is the distance from u to v (see Figure 5).

We now analyze $D\left(\boldsymbol{G}_s^u\right) + D\left(\boldsymbol{G}_s^v\right)$. By choice of s, we have

$$\begin{aligned} 2C\frac{s}{\log\log s} &= 2C\frac{\frac{n}{2} - \frac{1}{2}\alpha(n)}{\log\log\left(\frac{n}{2} - \frac{1}{2}\alpha(n)\right)} \\ &\overset{(\text{assume } \alpha(n)\le\frac{1}{2}n)}{\le} C\frac{n - \alpha(n)}{\log\log\left(\frac{n}{4}\right)} \\ &= C\frac{n\left(1 - \frac{\alpha(n)}{n}\right)}{\log\log n + \log\left(1 - \frac{\log 4}{\log n}\right)} \\ &\le C\frac{n\left(1 - \frac{\alpha(n)}{n}\right)}{\log\log n - \frac{2\log 4}{\log n}} = C\frac{n}{\log\log n}\frac{1 - \frac{\alpha(n)}{n}}{1 - \frac{2\log 4}{\log n \log\log n}} \\ &\le C\frac{n}{\log\log n}\left(1 - \frac{\alpha(n)}{n}\right)\left(1 + \frac{6}{\log n \log\log n}\right) \\ &\le C\frac{n}{\log\log n}\left(1 + \frac{6}{\log n \log\log n} - \frac{\alpha(n)}{n}\right) \end{aligned}$$

Since we assume that E_s^u and E_s^v hold, that is, that $D\left(\boldsymbol{G}_s^v\right), D\left(\boldsymbol{G}_s^u\right) \leq Cs\frac{\log\log\log s}{\log\log s} \leq Cs\frac{\log\log\log n}{\log\log s}$, we have

$$\begin{aligned} &D\left(\boldsymbol{G}_s^u\right) + D\left(\boldsymbol{G}_s^v\right) \\ &\quad \leq Cn\frac{\log\log\log n}{\log\log n}\left(1 + \frac{6}{\log n \log\log n} - \frac{\alpha(n)}{n}\right) \end{aligned}$$

Choosing $\alpha(n) = \frac{17n}{\log n}$ then gives

$$D\left(\boldsymbol{G}_s^u\right) + D\left(\boldsymbol{G}_s^v\right) \leq Cn\frac{\log\log\log n}{\log\log n}\left(1 - \frac{1}{\log n}\right)$$

Choosing also $\beta(n) = \frac{\log 2}{18}\log n \log\log n$, so that $\ell = \frac{18}{\log 2}\frac{n}{\log n \log\log n}$, by (5) we have that

$$\begin{aligned} d_{\boldsymbol{G}_n}(x,y) &\leq \frac{36}{\log 2}\frac{n}{\log n \log\log n} \\ &\qquad + Cn\frac{\log\log\log n}{\log\log n}\left(1 - \frac{1}{\log n}\right) + 1 \\ \text{(for } C \text{ large enough)} &\leq Cn\frac{\log\log\log n}{\log\log n} \end{aligned}$$

All that remains is to bound the probability of the event $\{U(x) \sim U(y) \forall x, y \in V_n\}$ from below. We do this using a union bound. The number of vertices in $S_\ell(x)$ can readily be seen to be

$$\sum_{k_1=s}^{n-1}\sum_{k_2=s}^{k_1-1}\sum_{k_3=s}^{k_2-1}\cdots\sum_{k_\ell=s}^{k_{\ell-1}-1} 1 = \sum_{k_1=1}^{n-s-1}\sum_{k_2=1}^{k_1-1}\sum_{k_3=1}^{k_2-1}\cdots\sum_{k_\ell=1}^{k_{\ell-1}-1} 1$$

This is the number of decreasing positive integer sequences of length ℓ, whose maximum entry is bounded by $n-s-1$. Since every choice of ℓ integers can be ordered in a unique fashion, we have

$$\begin{aligned} \left|S_\ell(x)\right| &= \binom{n-s-1}{\ell} \\ &= \binom{\frac{n}{2} + \frac{1}{2}\alpha(n) - 1}{n/\beta(n)} \\ &\geq \left(\frac{\frac{n}{2} + \frac{1}{2}\alpha(n) - 1}{n/\beta(n)}\right)^{n/\beta(n)} \\ &\geq \left(\frac{\beta(n)}{2}\right)^{n/\beta(n)} \\ &= \exp\left((\log\beta(n) - \log 2)\frac{n}{\beta(n)}\right) \\ &\quad \text{(Assume} n_0 \text{ large so that } \log\beta(n) \geq 2\log 2) \\ &\geq \exp\left(\frac{n\log\beta(n)}{2\beta(n)}\right) \end{aligned}$$

The collection $U(x) = \cup_{(z,t)\in S_\ell(x)} I_s(z)$ has size at least

$$|U(x)| = 2^s\left|S_\ell(x)\right| \geq 2^{\frac{n}{2} - \frac{1}{2}\alpha(n) + \frac{\log 2}{2}\frac{n\log\beta(n)}{\beta(n)}}$$

Denoting $U = |U(x)| = |U(y)|$, the probability that the sets $U(x)$ and $U(y)$ are disconnected at the nth step is therefore bounded from above by

$$\begin{aligned} &\mathbb{P}[U(x) \nsim U(y)] \\ &\quad = \frac{\binom{2^{n-1}-U}{U}}{\binom{2^{n-1}}{U}} \\ &\quad = \frac{\left(2^{n-1}-U\right)\left(2^{n-1}-U-1\right)\cdots\left(2^{n-1}-2U+1\right)}{2^{n-1}\left(2^{n-1}-1\right)\cdots\left(2^{n-1}-U+1\right)} \\ &\quad \text{(AM-GM inequality)} \\ &\quad \leq \frac{\left(2^{n-1} - \frac{3U-1}{2}\right)^U}{\left(2^{n-1}-U\right)^U} \\ &\quad = \left(1 - \frac{1}{2}\frac{U-1}{2^{n-1}-U}\right)^U \\ &\quad \leq \left(1 - \frac{U-1}{2^n}\right)^U \\ &\quad \leq \exp\left(-U^2/2^n + U/2^n\right) \\ &\quad \leq \exp\left(-2^{-\alpha(n) + \log 2\frac{n\log\beta(n)}{\beta(n)} + o(1)}\right) \end{aligned}$$

Plugging in our choice of $\alpha(n)$ and $\beta(n)$, we get

$$\begin{aligned} \mathbb{P}[U(x) \nsim U(y)] &\leq \exp\left(-2^{-\frac{17n}{\log n} + \log 2\frac{n\log\left(\frac{\log 2}{18}\log n\log\log n\right)}{\frac{\log 2}{18}\log n \log\log n} + o(1)}\right) \\ &\leq \exp\left(-2^{-\frac{17n}{\log n} + \frac{18n}{\log n}(1+o(1))}\right) \\ &= \exp\left(-2^{\frac{n}{\log n}(1+o(1))}\right) \\ &\leq \exp\left(-2^{\frac{n}{2\log n}}\right) \end{aligned}$$

for n large enough. As there are no more than $2^{2n} = e^{2n\log 2}$ choices for the pairs x, y, this gives

$$\begin{aligned} \mathbb{P}[\exists x, y \text{ s.t. } U(x) \nsim U(y)] &\leq \exp\left(-2^{\frac{n}{2\log n}} + 2n\log 2\right) \\ &\leq \exp\left(-2^{\frac{n}{4\log n}}\right) \end{aligned}$$

for n large enough. Let $E_k = \cup_{z\in\{0,1\}^n} E_k^z$. We have thus shown that

$$\mathbb{P}[E_n | E_1, \ldots, E_{n-1}] \geq 1 - \exp\left(-2^{\frac{n}{4\log n}}\right)$$

Since there are 2^{n-k} instances of $\boldsymbol{G}_k$ in $\boldsymbol{G}_n$, and recalling that $\mathbb{P}[E_k] = 1$ for $k \leq k^*$, by repeated conditioning we thus have

$$\begin{aligned} \mathbb{P}[E_n] &\geq 1 - \sum_{k=k^*}^{n}\exp\left(-2^{\frac{k}{4\log k}}\right)2^{n-k} \\ &\geq 1 - \sum_{k=k^*}^{n}\exp\left(-2^{\frac{k}{4\log k}} + n\log 2\right) \end{aligned}$$

By choosing γ large enough, for $k \geq k^*$ we have

$$\frac{k}{4\log k} \geq \frac{k^*}{4\log k^*} = \frac{\gamma \log n \log\log n}{4\log(\gamma \log n \log\log n)} \geq 2\log_2 n$$

and so

$$\mathbb{P}[E_n] \geq 1 - \sum_{k=k^*}^{n} e^{-n}$$

4.3 | Vertex Expansion

The proof of Theorem 3 relies on the observation that a set $S \subseteq V_n$ sampled uniformly at random will be an α-vertex expander with high probability (for some small constant $\alpha > 0$), since a constant fraction of the edges of σ_{n-1} will go from S to its complement. This alone is not enough, since there are always sets of the form $S = S_0 \cup S_1$, where $S_0 \subseteq V_n^0$ (recall (4) for the definition of V_n^z) and $S_1 = \{N_n(x) | x \in S_0\}$. To overcome this, we look at edges coming from the last three permutations, $\sigma_{n-1}, \sigma_{n-2}, \sigma_{n-3}$, and bound the number of sets $S \subseteq V_n$ so that the boundary that comes from σ_{n-3}-edges isn't large enough. Afterwards we apply a union bound over these sets to bound the probability that they have a small σ_{n-1}- and σ_{n-2}-boundary.

More precisely, sets which have a small contribution to their boundary at the kth generation are defined as follows.

Definition 4. (Badly-matched sets). Let $x \in \{0,1\}^k$. Let $A \subseteq V_k^0, B \subseteq V_k^1$. We say that A, B are (k,α)-*badly-matched* if

$$\frac{2\left|x \in A | N_k(x) \in B\right|}{|A| + |B|} \geq (1-\alpha)$$

Remark 8. If A, B are badly-matched, then $||A| - |B|| \leq \alpha(|A| + |B|)$. This is because if, say, $|A| > |B| + \alpha(|A| + |B|)$ then even when all edges from B go into A there will still be $\alpha(|A| + |B|)$ edges between A and $V_k^1 \setminus B$. This implies that

$$2\left|x \in A | N_k(x) \in B\right| \leq 2|B| < (1-\alpha)(|A| + |B|)$$

If A, B are not badly-matched, then $\left|\partial_k(A \cup B)\right| > \alpha|A \cup B|$, since $\left|\partial_k(A \cup B)\right| = |A| + |B| - 2\left|x \in A | N_k(x) \in B\right|$, so the set $A \cup B$ has α-expansion. If A, B are (k,α)-badly-matched, then they are also (k,α')-badly-matched for every $\alpha' \geq \alpha$.

As alluded to above, we start by bounding the possible number of badly-matched sets in generation $n-2$, for any permutation σ_{n-3}; this is the content of Proposition 4. We then bound the probability that said badly-matched sets are also badly-matched in generations $n-1$ and n; this is the content of Proposition 5. The last claim we need for the proof is that sets of size $O(n)$ have non-trivial vertex expansion regardless of the permutation. The proofs of all assertions are found at the end of the section.

Recall that for an integer $k > 0$, we set $K := 2^k$. In addition, denote by $H(x) = -x\log x - (1-x)\log(1-x)$ the binary entropy function.

Proposition 4. *There exists a function $\delta : \mathbb{R} \to \mathbb{R}$ with $\lim_{x\to 0}\delta(x) = 0$, that depends on η, such that the following holds. Let $0 < \alpha \leq \frac{1}{2}$, and let $k, j > 0$ be integers so that $j \leq \eta K$. For any permutation σ_{k-1}, the number of (k,α)-badly-matched sets $A \subseteq V_k^0$ and $B \subseteq V_k^1$ such that $(1-\alpha)\frac{j}{2} \leq |A|, |B| \leq (1+\alpha)\frac{j}{2}$ is smaller than*

$$5\alpha^3 K^3 2^{\frac{K}{2}(1+\delta(\alpha))H(\frac{j}{K}) + j\delta(\alpha)}$$

Proposition 5. *There exists a function $\delta : \mathbb{R} \to \mathbb{R}$ with $\lim_{x\to 0}\delta(x) = 0$ that depends on η, such that the following holds. Let $0 < \alpha \leq \frac{1}{2}$, and let $k, j > 0$ be integers so that $j \leq \eta K$. Let $A \subseteq V_k^0$ and $B \subseteq V_k^1$ be such that $(1-\alpha)\frac{j}{2} \leq |A|, |B| \leq (1+\alpha)\frac{j}{2}$. If the permutations $\overline{\sigma}$ are uniformly random, then*

$$\mathbb{P}[A, B \text{ are } (k,\alpha)\text{-badly-matched}] \leq 3\alpha K^2 2^{-\frac{K}{2}(1-\delta(\alpha))H(\frac{j}{K}) + \delta(\alpha) j} \tag{6}$$

Claim 1. Let $c > 3$. Then there is some $n_0 \in \mathbb{N}$ so that for every $n > n_0$ and every $S \subseteq V_n$ so that $|S| \leq cn$,

$$|\partial S| \geq \frac{1}{c^2}|S|$$

Proof of Theorem 3. Fix $\eta > 0$ and let $\alpha \in (0, 1/2]$ be chosen later. We define F_n to be the event that $\boldsymbol{G}_n$ is not an (η,α)-vertex expander. We will show that

$$\lim_{n\to\infty} \mathbb{P}[F_n] = 0$$

We will bound

$$\mathbb{P}[F_n] \leq \sum_{j=0}^{\eta N} \mathbb{P}[\exists S \text{ s.t. } |S| = j \text{ and } |\partial S| < \alpha|S|] \tag{7}$$

By Claim 1, it is enough to start this sum with $j = cn$, as long as $\alpha \leq \frac{1}{c^2}$. The constant c will be determined at the end of the proof. Thus the right-hand side of (7) is equal to

$$\sum_{j=cn}^{\eta N} \mathbb{P}[\exists S \text{ s.t. } |S| = j \text{ and } |\partial S| < \alpha|S|]$$

Let $S \subseteq V_n$ with $|S| = j$. For $x \in \{0,1\}^k$, denote $S_x = S \cap V_n^x$. If the sets S_0, S_1 are not (n,α)-badly-matched, then by Remark 8, the edges from σ_{n-1} are enough to guarantee a large boundary, that is, the set S has α-expansion. This happens in particular when $\left|\left|S_1\right| - \left|S_0\right|\right| \geq \alpha|S|$. We may thus restrict ourselves to S that satisfy $\left|S_i\right| \geq \frac{1}{2}(1-\alpha)|S|$. Thus

$$(1-\alpha)\frac{j}{2} \leq \left|S_i\right| \leq (1+\alpha)\frac{j}{2} \tag{8}$$

for every $i \in \{0,1\}$. Similarly, if S_{i0}, S_{i1} are not $(n-1, 3\alpha)$-badly-matched for any $i \in \{0,1\}$, then $\left|\partial_{n-1}(S_{i0} \cup S_{i1})\right| \geq 3\alpha\left|S_i\right| \geq \frac{3}{2}\alpha(1-\alpha)|S|$, which is larger than $\alpha|S|$ for α small enough. This happens in particular when $\left|\left|S_{i1}\right| - \left|S_{i0}\right|\right| \geq 3\alpha\left|S_i\right|$. We may thus further restrict ourselves to S that satisfy $\left|S_{ij}\right| \geq \frac{1}{2}(1-3\alpha)\left|S_i\right|$, which means that

$$(1-3\alpha)(1-\alpha)\frac{j}{4} \leq \left|S_{ij}\right| \leq (1+3\alpha)(1+\alpha)\frac{j}{4} \tag{9}$$

for every $i, j \in \{0,1\}$.

Finally, if there is are sets S_{ij0}, S_{ij1} for some $i, j \in \{0,1\}$ that are not $(n-2, 5\alpha)$-badly-matched, then for α small enough we have $\left|\partial_{n-2}(S_{ij0} \cup S_{ij1})\right| \geq \alpha|S|$, and we can assume that

$$(1-5\alpha)(1-3\alpha)(1-\alpha)\frac{j}{8} \leq \left|S_{ij}\right| \leq (1+5\alpha)(1+3\alpha)(1+\alpha)\frac{j}{8}$$

In particular, this happens when

$$(1-10\alpha)\frac{j}{8} \leq \left|S_{ij}\right| \leq (1+10\alpha)\frac{j}{8}$$

Thus, to bound $\mathbb{P}[\exists S \text{ s.t. } |S| = j \text{ and } |\partial S| < \alpha|S|]$, we only need to consider sets S whose all four pairs S_{ij0}, S_{ij1} are $(n-2, 10\alpha)$ badly-matched; any other set has α-expansion by Remark 8. By Proposition 4, with $k = n-2$, for each $i, j \in \{0,1\}$ there are at most $5000\alpha^3 N^3 2^{\frac{N}{8}(1+\delta(10\alpha))H\left(\frac{j}{N}\right)+j\delta(10\alpha)}$ sets S_{ij} such that S_{ij0}, S_{ij1} are $(n-2, 10\alpha)$-badly-matched, so there are at most

$$\left(5000\alpha^3 N^3 2^{\frac{N}{8}(1+\delta(10\alpha))H\left(\frac{j}{N}\right)+j\delta(10\alpha)}\right)^4 = 5000^4\alpha^{12}N^{12}2^{\frac{N}{2}(1+\delta(10\alpha))H\left(\frac{j}{N}\right)+4j\delta(10\alpha)}$$

possible sets to consider. Thus

$$\begin{aligned}&\mathbb{P}[\exists S \text{ s.t. } |S| = j \text{ and } |\partial S| < \alpha|S|]\\ &\leq 5000^4\alpha^{12}N^{12}2^{\frac{N}{2}(1+\delta(10\alpha))H\left(\frac{j}{N}\right)+4j\delta(10\alpha)}\cdot\\ &\quad\cdot \max_S \mathbb{P}[S \text{ does not have } \alpha\text{-expansion}]\end{aligned} \tag{10}$$

where S is restricted as above. To bound the probability, observe that for any fixed $S \subseteq V_n$,

$$\begin{aligned}&\mathbb{P}[S \text{ does not have } \alpha\text{-expansion}]\\ &\leq \mathbb{P}[S_0, S_1 \text{ are } (n,\alpha)\text{-badly-matched}\\ &\quad \text{and } S_{00}, S_{01} \text{ are } (n-1, 3\alpha)\text{-badly-matched}]\end{aligned}$$

As the event $\{S_0, S_1$ are badly-matched$\}$ depends only on the permutation σ_{n-1} and $\{S_{00}, S_{11}$ are badly-matched$\}$ depends only on the permutation σ_{n-2}, these two events are independent. By the relations (8) and (9), we can apply Proposition 5, yielding

$$\begin{aligned}&\mathbb{P}[S \text{ does not have } \alpha\text{-expansion}]\\ &\leq \mathbb{P}[S_0, S_1 \text{ are not } (n,\alpha)\text{-badly-matched}]\\ &\quad \cdot \mathbb{P}[S_{00}, S_{01} \text{ are not } (n-1, 3\alpha)\text{-badly-matched}]\\ &\leq 3\alpha N^2 2^{-\frac{N}{2}(1-\delta(\alpha))H\left(\frac{j}{N}\right)+\delta(\alpha)j} \cdot 3\alpha N^2 2^{-\frac{N}{4}(1-\delta(3\alpha))H\left(\frac{j}{N}\right)+\delta(3\alpha)j}\end{aligned}$$

For simplicity, in the next inequalities we unify all the expressions of the form $\delta(c\alpha)$ appearing in the exponents to $\delta(\alpha)$ (that goes to 0 as $\alpha \to 0$). Using (10), we get that,

$$\begin{aligned}&\mathbb{P}[\exists S \text{ s.t. } |S| = j \text{ and } |\partial S| < \alpha|S|]\\ &\leq 9 \times 5000^4\alpha^{14}N^{16}2^{\delta(\alpha)j}(2^{\frac{1}{4}(1-\delta(\alpha))})^{-NH\left(\frac{j}{N}\right)}\end{aligned}$$

Note that we abused notation. Plugging this back in (7) we get

$$\mathbb{P}[F_n] \leq 9 \times 5000^4\alpha^{14}N^{16}\sum_{j=cn}^{\eta N} 2^{\delta(\alpha)j}\left(2^{\frac{1}{4}(1-\delta(\alpha))}\right)^{-NH\left(\frac{j}{N}\right)} \tag{11}$$

We take α so that $\delta(\alpha) < \frac{1}{2}$ and get that $2^{\frac{1}{4}(1-\delta(\alpha))} \geq 2^{\frac{1}{8}}$. In addition, we use the well known inequality $H(x) \geq 4x(1-x)$ to bound the right-hand side of (11) from above by

$$\begin{aligned}&9 \times 5000^4\alpha^{14}N^{16}\sum_{j=cn}^{\eta N} 2^{\delta(\alpha)j}2^{-\frac{j}{2}(1-\eta)}\\ &= 9 \times 5000^4\alpha^{14}N^{16}\sum_{j=cn}^{\eta N} 2^{-\frac{j}{2}(1-\eta-2\delta(\alpha))}\end{aligned}$$

Finally, by taking α so that $1-\eta-2\delta(\alpha) \geq \frac{1-\eta}{2}$ and taking c so that $2^{\frac{cn/2}{2}(\frac{1-\eta}{2})} \geq N^{16} = 2^{16n}$ we get that the sum on the right-hand side is at most

$$\begin{aligned}&9 \times 5000^4\alpha^{14}N^{16}2^{-\frac{cn/2}{2}(\frac{1-\eta}{2})}\sum_{j=\frac{c}{2}n}^{\infty} 2^{-\frac{j}{4}(1-\eta)}\\ &\leq 9 \times 5000^4\alpha^{14}\sum_{j=\frac{c}{2}n}^{\infty} 2^{-\frac{j}{4}(1-\eta)}\\ &= \frac{9}{1-2^{-(1-\eta)/4}} \cdot 5000^4\alpha^{14}2^{-\frac{cn}{8}(1-\eta)}\end{aligned}$$

Thus,

$$\mathbb{P}[F_n] \leq \frac{9}{1-2^{-(1-\eta)/4}} \cdot 5000^4\alpha^{14}2^{-\frac{cn}{8}(1-\eta)}$$

This tends to 0 as $n \to \infty$.

Proof of Proposition 4. We can count the subsets $A \subseteq V_k^0, B \subseteq V_k^1$ by first choosing a set of edges of the kth matching that are connected to $A \cup B$. For each chosen edge $\{x, y\}$ where $x \in V_k^0, y \in V_k^1$ we decide whether $x \in A, y \in B$, or $x \in A, y \notin B$ or $x \notin A, y \in B$. For sets that are (k,α)-badly matched, our count yields the following.

1. The number of edges that are adjacent to $A \cup B$ is at least $(1-\alpha)\frac{j}{2}$ (the lower bound is achieved when $|A| = |B| = (1-\alpha)\frac{j}{2}$ and $N_k(A) = B$). It is at most $(1+\alpha)\frac{j}{2} + 2\alpha(1+\alpha)\frac{j}{2}$ (since there could be at most $(1+\alpha)\frac{j}{2}$ that cross from A to B, and no more than $2\alpha(1+\alpha)\frac{j}{2}$ additional edges that are adjacent to only one of A, B, since A, B are supposed to be (k,α)-badly matched). Since $(1+\alpha)\frac{j}{2} + 2\alpha(1+\alpha)\frac{j}{2} \leq (1+4\alpha)\frac{j}{2}$ for $\alpha \leq \frac{1}{2}$, using the relation $\binom{n}{k} \leq 2^{nH(k/n)}$, the number of possible choices for edges adjacent to $A \cup B$ is at most

$$\begin{aligned}\sum_{\ell=(1-\alpha)\frac{j}{2}}^{(1+4\alpha)\frac{j}{2}} \binom{\frac{K}{2}}{\ell} &\leq \sum_{\ell=(1-\alpha)\frac{j}{2}}^{(1+4\alpha)\frac{j}{2}} 2^{\frac{K}{2}H(\frac{\ell}{K/2})}\\ &\leq \sum_{\ell=(1-\alpha)\frac{j}{2}}^{(1+4\alpha)\frac{j}{2}} 2^{\frac{K}{2}H\left(\frac{j}{K}\right)+\frac{K}{2}\left(H\left(\frac{\ell}{K/2}\right)-H\left(\frac{j/2}{K/2}\right)\right)}\end{aligned} \tag{12}$$

By Lagrange's mean-value theorem, $\frac{K}{2}(H\left(\frac{\ell}{K/2}\right) - H\left(\frac{j/2}{K/2}\right)) = \left(\ell - \frac{j}{2}\right)H'(\xi) = -\left(\ell - \frac{j}{2}\right)\log\frac{\xi}{1-\xi}$ for some

ξ between $\frac{\ell}{K/2}$ and $\frac{j/2}{K/2}$. Thus we can write $\xi = c' \frac{j/2}{K/2}$ for some $(1-\alpha) \le c' \le (1+4\alpha)$, and we have

$$\left| \left(\ell - \frac{j}{2} \right) H'(\xi) \right| \le 4\alpha j |H'(\xi)| = 4\alpha j |\log \xi - \log(1-\xi)| \tag{13}$$

We now bound the logarithms. Since $\xi = c' \frac{j}{K} \le c'\eta$, we have that $-\log(1-\xi) \le -\log(1-c'\eta)$; the quantity on the right-hand side is just a constant (provided that α is small enough so that $(1+4\alpha)\eta < 1$). For $|\log \xi|$, we have

$$\begin{aligned} 4\alpha j |\log \xi| &= 4\alpha j \left| \log\left(c' \frac{j}{K} \right) \right| \le 4\alpha j |\log c'| + 4\alpha j \log \frac{j}{K} \\ &= 4\alpha j |\log c'| + 8\frac{K}{2}\alpha \frac{j}{K} \log \frac{j}{K} \\ &\le 4\alpha j |\log c'| + 8\frac{K}{2}\alpha H\left(\frac{j}{K} \right) \end{aligned}$$

Thus, there exists a constant $c > 0$ (that depends on η) such that

$$\left| \left(\ell - \frac{j}{2} \right) H'(\xi) \right| \le c\alpha j + c\alpha \frac{K}{2} H\left(\frac{j}{K} \right)$$

for some $c > 0$ that depends on η. Thus the left-hand side in (12) is at most

$$5\alpha K 2^{(1+c\alpha)\frac{K}{2}H(\frac{j}{K})+c\alpha j} \tag{14}$$

2. Then we choose out of the edges adjacent to $A \cup B$ the edges that touch A only, and the edges that touch B only. As A, B are (k, α)-badly matched, at least a $(1-\alpha)$-fraction of the edges must touch both A and B, so no more than an α-fraction of the edges are available to touch only one of the sets. Assuming that $\alpha < 1/4$, the number of possibilities (for a given edge set chosen in the previous step) is at most

$$\begin{aligned} &\sum_{\ell_A=0}^{\alpha(1+4\alpha)\frac{j}{2}} \binom{(1+4\alpha)\frac{j}{2}}{\ell_A} \sum_{\ell_B=0}^{\alpha(1+4\alpha)\frac{j}{2}} \binom{(1+4\alpha)\frac{j}{2}}{\ell_B} \\ &\le (\alpha j)^2 \binom{(1+4\alpha)\frac{j}{2}}{\alpha(1+4\alpha)\frac{j}{2}}^2 \le \alpha^2 K^2 2^{2jH(\alpha)} \end{aligned} \tag{15}$$

Multiplying (14) and (15), and setting $\delta(\alpha) = 2H(\alpha) + c\alpha$, the number of badly-matched pairs is bounded by

$$5\alpha^3 K^3 2^{\frac{K}{2}(1+\delta(\alpha))H(\frac{j}{K})+\delta(\alpha)j}$$

□

Proof of Proposition 5. To bound the probability in (6), we go over all possible subsets $A' \subseteq A$ and sum the probability that the set of outgoing edges from A' is some set $B' \subseteq B$. Since A and B both have sizes in the interval $\left[(1-\alpha)\frac{j}{2}, (1+\alpha)\frac{j}{2}\right]$, the size of A', B' should be at least $(1-\alpha)(1-\alpha)\frac{j}{2} \ge (1-2\alpha)\frac{j}{2}$. The probability is bounded by

$$\begin{aligned} &\sum_{\ell=(1-2\alpha)\frac{j}{2}}^{(1+\alpha)\frac{j}{2}} \sum_{\substack{A' \subseteq A, B' \subseteq B \\ |A'|=|B'|=\ell}} \mathbb{P}\left[N_k(A') = B'\right] \\ &= \sum_{\ell=(1-2\alpha)\frac{j}{2}}^{(1+\alpha)\frac{j}{2}} \sum_{\substack{A' \subseteq A, B' \subseteq B \\ |A'|=|B'|=\ell}} \frac{1}{\binom{\frac{K}{2}}{\ell}} \\ &\le \sum_{\ell=(1-2\alpha)\frac{j}{2}}^{(1+\alpha)\frac{j}{2}} \sum_{\substack{A' \subseteq A, B' \subseteq B \\ |A'|=|B'|=\ell}} \frac{K}{2} 2^{-\frac{K}{2}H(\frac{\ell}{K/2})} \\ &\overset{\left(\text{assuming } \alpha < \frac{1}{5}\right)}{\le} \sum_{\ell=(1-2\alpha)\frac{j}{2}}^{(1+\alpha)\frac{j}{2}} \binom{(1+\alpha)\frac{j}{2}}{(1-2\alpha)\frac{j}{2}}^2 \frac{K}{2} 2^{-\frac{K}{2}H(\frac{\ell}{K/2})} \\ &\le 2^{2(1+\alpha)\frac{j}{2}H(\frac{1-2\alpha}{1+\alpha})} \frac{K}{2} \sum_{\ell=(1-2\alpha)\frac{j}{2}}^{(1+\alpha)\frac{j}{2}} 2^{-\frac{K}{2}H(\frac{\ell}{K/2})} \end{aligned} \tag{16}$$

By Lagrange's mean-value theorem, we write

$$\begin{aligned} \frac{K}{2}H\left(\frac{\ell}{K/2} \right) &= \frac{K}{2}H\left(\frac{j}{K} \right) + \frac{K}{2}\left(H\left(\frac{\ell}{K/2} \right) \right. \\ &\left. -H\left(\frac{j/2}{K/2} \right) \right) = \frac{K}{2}H\left(\frac{j}{K} \right) + \left(\ell - \frac{j}{2} \right) H'(\xi) \end{aligned}$$

for some ξ between $\frac{\ell}{K/2}$ and $\frac{j/2}{K/2}$. As $|\ell - \frac{j}{2}| \le \alpha j$, we bound $(\ell - \frac{j}{2})H'(\xi)$ by $\alpha j |H'(\xi)|$. Write $\xi = c' \frac{j}{K}$ for some $1 - 2\alpha \le c' \le 1 + 2\alpha$. Then (similar to (13) in the proof of the previous proposition)

$$\left| (\ell - \frac{j}{2}) H'(\xi) \right| \le c\alpha j + c\alpha \frac{K}{2} H\left(\frac{j}{K} \right)$$

for some constant $c > 0$ which only depends on η. Thus (16) is at most

$$\begin{aligned} &2^{2(1+\alpha)\frac{j}{2}H\left(\frac{1-2\alpha}{1+\alpha}\right)} \frac{K}{2} \sum_{\ell=(1-2\alpha)\frac{j}{2}}^{(1+\alpha)\frac{j}{2}} 2^{-\frac{K}{2}(1-\delta(\alpha))H(\frac{j}{K})+\delta(\alpha)j} \\ &\le 3\alpha K^2 2^{-\frac{K}{2}(1-\delta(\alpha))H(\frac{j}{K})+\delta(\alpha)j} \end{aligned}$$

where $\delta(\alpha) = \max\left\{ (1+\alpha)H\left(\frac{1-2\alpha}{1+\alpha} \right), c\alpha \right\}$. □

Proof of Claim 1. For every vertex $v \in V_n$, the second neighborhood of v, $A_n(v) := \{u \in V_n | d(v,u) = 2\}$, is of size at least $\binom{n}{2}$. This can be seen by induction. The base case for $n = 2$ is clear. Assume without loss of generality that $v \in V_n^0$ and partition $A_n(v) = \left(A_n(v) \cap V_n^0\right) \cup \left(A_n(v) \cap V_n^1\right)$. Note that in the instance of G_{n-1} whose vertex set is V_0, the second neighborhood of v is $A_{n-1}(v) = A_n(v) \cap V_n^0$. Thus, by the induction hypothesis, $|A_n(v) \cap V_n^0| \ge \binom{n-1}{2}$. In addition, $A_n(v) \cap V_n^1$ contains the neighborhood of $N_n(v)$ inside V_n^1, which is of size $n-1$. Summing up sizes we get $|A_n(v)| \ge \binom{n-1}{2} + \binom{n-1}{1} = \binom{n}{2}$. Note that for the hypercube Q_n we have strict equality.

Now fix $S \subseteq V_n$ of size at most cn and let $v \in S$. If a $\frac{1}{c}$-fraction of the neighborhood of v is not in S then $|\partial S| \ge \frac{1}{c}n \ge \frac{1}{c^2}|S|$. Otherwise, at least $(1 - \frac{1}{c})$-fraction of v's neighbors are inside S. Denote these vertices as $T := N(v) \cap S$. Thus $|\partial S| \geq |N(T)| - |S|$. Since the neighborhood of v is $A_n(v) \cup \{v\}$, the neighborhood of T is of size at least

$$|A_n(v)| - n|N(v) \setminus S| \geq \binom{n}{2} - \frac{1}{c}n^2 \geq \frac{1}{7}n^2$$

Hence $|\partial S| \geq \frac{1}{7}n^2 - cn \geq \frac{1}{c^2}|S|$ for a large enough n. □

4.4 | Eigenvalues

Proof of Proposition 3. Since G_n is n-regular, its largest eigenvalue λ_1 is n, and its corresponding eigenvector $f_n : \{0,1\}^n \to \mathbb{R}$ satisfies $f_n(x) = 1$. To show that $\lambda_2 \geq n-2$, let $g_n : \{0,1\}^n \to \mathbb{R}$ be given by

$$g_n(x) = (-1)^{x_n}$$

that is, g_n takes value 1 on the first instance of G_{n-1} in G_n, and -1 on the second instance. Then $A_n g_n = (n-2)g_n$.

The proof that $\lambda_2 \leq n-2$ is by induction. The claim clearly holds for $n = 1$, where G_1 is just an edge. Assume it holds for all $k \leq n-1$, and let $h : \{0,1\}^n \to \mathbb{R}$ be an eigenvector of A_n that is orthogonal to both f_n and g_n. If we write

$$h(x) = \begin{cases} h_0(x_1, \dots, x_{n-1}) & x_n = 0 \\ h_1(x_1, \dots, x_{n-1}) & x_n = 1 \end{cases}$$

for some functions $h_i : \{0,1\}^{n-1} \to \mathbb{R}$, then both h_0 and h_1 are orthogonal to f_{n-1}, and by the induction hypothesis, we have $h_i^T A_{n-1} h_i \leq (n-3)\|h_i\|_2^2$. Using the recursive matrix representation (1) of the twisted hypercube graph, we can write

$$h^T A_n h = \begin{pmatrix} h_0^T & h_1^T \end{pmatrix} \begin{pmatrix} A_{n-1}^0 & P \\ P^T & A_{n-1}^1 \end{pmatrix} \begin{pmatrix} h_0 \\ h_1 \end{pmatrix}$$

where P is the $2^{n-1} \times 2^{n-1}$ permutation matrix representing σ_{n-1}, and A_{n-1}^0 and A_{n-1}^1 are the adjacency matrices of the two instances of G_{n-1}. Explicitly opening the products, we get

$$\begin{aligned} h^T A_n h &= \begin{pmatrix} h_0^T & h_1^T \end{pmatrix} \begin{pmatrix} A_{n-1}^0 h_0 + P h_1 \\ P^T h_0 + A_{n-1}^1 h_1 \end{pmatrix} \\ &= h_0^T A_{n-1}^0 h_0 + h_0^T P h_1 + h_1^T P^T h_0 + h_1^T A_{n-1}^1 h_1 \\ &\leq (n-3)\|h_0\|_2^2 + 2h_0^T P h_1 + (n-3)\|h_1\|_2^2 \\ &= (n-3)\|h\|_2^2 + 2h_0^T P h_1 \leq (n-3)\|h\|_2^2 + 2\|h_0\|_2 \|h_1\|_2 \\ &\leq (n-3)\|h\|_2^2 + \|h_0\|_2^2 + \|h_1\|_2^2 = (n-2)\|h\|_2^2 \end{aligned}$$

□

Proof of Lemma 1. In the following, C is a constant depending on k whose value may change from instance to instance. A set of edges $F \subseteq E(\boldsymbol{G}_n)$ is said to be "finalized at generation m" if for every edge $\{x, y\} \in F$, $\gamma(x, y) \leq m$, and there exists at least one edge such that $\gamma(x, y) = m$. For a given $u \in V_n$, let $w = N_m(u)$ be its m-neighbor, and let $E_m(u)$ be event that there exists a cycle of length no more than k which contains the edge $\{u, w\}$ and is finalized at generation m.

We will now bound the probability of the event $E_m(u)$. Since $I_{m-1}(u) \neq I_{m-1}(w)$, that is, u and w are found on different copies of V_{m-1}, in order for a cycle of length $\leq k$ to exist, there must also be an m-generation edge going from $I_{m-1}(w)$ back to $B_{<m}(u, k)$; otherwise, any path starting with the edge $\{u, w\}$ cannot reach u again. In fact, this edge must be reachable from w in at most k steps. Let W be the set of all $z \in I_{m-1}(w)$ such that there exists a simple path $P = (x_1, \dots, x_t)$ with the following properties:

1. P is a shortest path from w to z, and $t \leq k$.
2. $\gamma(x_i, x_{i+1}) \leq m$ for all $i = 1, \dots, t-1$.
3. $\gamma(x_{t-1}, x_t) < m$.
4. P does not contain the edge $\{u, w\}$.

In other words, W is the set of all vertices in $I_{m-1}(w)$ which can be reached from w by a path of at most k edges of generation at most m, and which can still send out an m-generation edge without backtracking. If there are no edges from W to $B_{<m}(u, k)$, then there is no cycle of length $\leq k$ which contains $\{u, w\}$ (see Figure 6 for a graphical depiction).

Given $z \in W$, the probability that $N_m(z) \in B_{<m}(u, k)$ depends only on the m-generation edges used in the path P. Since σ_{m-1} is uniform, we can bound this probability by

$$\mathbb{P}\left[N_m(z) \in B_{<m}(u, k) | z \in W\right] \leq \frac{|B_{<m}(u, k)|}{\max\left\{1, 2^{m-1} - k\right\}} \leq C\frac{m^{k+1}}{2^m}$$

for some $C > 0$ which depends on k (we subtract k in the denominator, since in the worst case the path from w to z has at most k m-generation edges from $I_{m-1}(w)$ to $I_{m-1}(u) \setminus B_{<m}(u, k)$). Since there are at most m^{k+1} vertices in W, taking the union bound gives

$$\mathbb{P}[E_m(u)] \leq C\frac{m^{2k+2}}{2^m}$$

Letting $E_m = \cup_{u \in B(v, 2k)} E_m(u)$, we then have

$$\mathbb{P}[E_m] \leq C\frac{m^{2k+2}}{2^m} n^{2k+1}$$

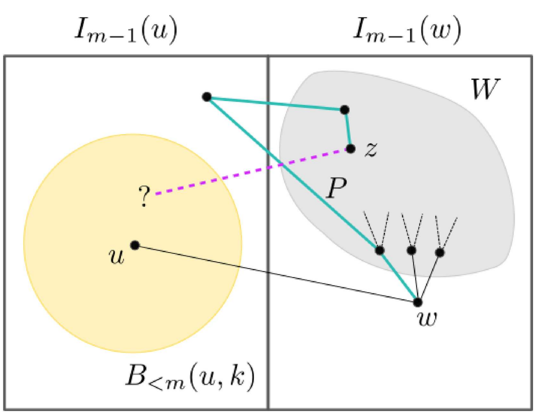

FIGURE 6 | There can be a cycle containing the edge $\{u, w\}$ only if there is an m-generation edge crossing from some $z \in W$ to $B_{<m}(u, k)$. Since both $B_{<m}(u, k)$ and W are small in comparison to $I_{m-1}(u)$, the probability of this happening is small.

In particular, there exists a constant $C > 0$ such that

$$\sum_{m>m_0} \mathbb{P}[E_m] \leq C2^{-m_0} m_0^{2k+2} n^{2k+1}$$

If a vertex $z \in B(v,k)$ is part of a cycle of length at most k which is finalized at generation m, then necessarily there exists some $u \in B(v,2k)$ such that $E_m(u)$ holds. Thus, if E_m^c holds for every $m > m_0$, then z can only be contained in cycles of length at most k which are finalized at generation $\leq m_0$. The number of such cycles is bounded by

$$\sum_{m=1}^{m_0} \sum_{i=1}^{k} m^i \leq \sum_{m=1}^{m_0} k m^k \leq C m_0^{k+1}$$

for some constant $C > 0$. The probability of F_v is then lower bounded by

$$\begin{aligned}\mathbb{P}[F_v] &\geq \mathbb{P}\left[\cap_{m>m_0} E_m^c\right] \\ &= 1 - \mathbb{P}\left[\cup_{m>m_0} E_m\right] \\ &\geq 1 - \sum_{m>m_0} \mathbb{P}[E_m] \\ &\geq 1 - C2^{-m_0} m_0^{2k+2} n^{2k+1}\end{aligned}$$

as needed.

Proof of Theorem 4. We use the moment method. While the main technique is classical (see e.g., [17]), we write the proof in full for completeness.

Proving that μ_n converges weakly to μ_{circ} in probability means that for every continuous function $f : \mathbb{R} \to \mathbb{R}$, we have convergence in probability of the expected value of f:

$$\int_{\mathbb{R}} f\, d\mu_n \xrightarrow{P} \int_{\mathbb{R}} f\, d\mu_{\text{circ}} \tag{17}$$

as $n \to \infty$. By the Weierstrass theorem, every continuous function on a closed interval can be arbitrarily well-approximated by a finite-degree polynomial. Since μ_{circ} is supported on a bounded interval, it suffices to show (17) for functions of the form $f_k = x^k$, that is, showing that the kth moments of μ_n converge the to kth moments of μ_{circ}. These moments are known, and are given by

$$\int_{\mathbb{R}} x^k d\mu_{\text{circ}} = \begin{cases} C_{k/2} & k \text{ is even} \\ 0 & k \text{ is odd} \end{cases}$$

where C_m is the mth Catalan number, and is equal to the number of ordered rooted trees with m edges. We will first show that $\mathbb{E} \int_{\mathbb{R}} x^k d\mu_n \to \int_{\mathbb{R}} x^k d\mu_{\text{circ}}$, and then show that $\mathrm{Var}\left(\int_{\mathbb{R}} x^k d\mu_n\right) \to 0$; by Chebyshev's inequality, this implies the desired convergence in probability.

Since μ_n is just the empirical measure of the eigenvalues of $A/\sqrt{n}$, we have

$$\begin{aligned}\int_{\mathbb{R}} x^k d\mu_n &= \frac{1}{2^n} \sum_{i=1}^{2^n} \left(\frac{\lambda_i}{\sqrt{n}}\right)^k = \frac{1}{2^n} \mathrm{Tr}\left(\frac{A}{\sqrt{n}}\right)^k \\ &= \frac{1}{2^n} \frac{1}{n^{k/2}} \sum_{i_1,\ldots,i_k=1}^{2^n} A_{i_1 i_2} A_{i_2 i_3} \cdots A_{i_{k-1} i_k} A_{i_k i_1}\end{aligned}$$

For a fixed i_1, the sum $\sum_{i_2,\ldots,i_k} A_{i_1 i_2} \cdots A_{i_k i_1}$ is the number of walks of length k in $\boldsymbol{G}_n$ that start and end at the vertex i_1. Let $X_v(t)$ be the simple random walk on $\boldsymbol{G}_n$ which starts at vertex v. Then, since $\boldsymbol{G}_n$ is n-regular, the number of simple random walks of length k is n^k, and we have

$$\int_{\mathbb{R}} x^k d\mu_n = \frac{1}{2^n} n^{k/2} \sum_{i=1}^{2^n} \mathbb{P}[X_i(k) = i] \tag{18}$$

where the probability is over the randomness induced by the random walk. Taking expectations over the measure induced by the permutations, we thus have, for any $v \in V_n$,

$$\begin{aligned}\mathbb{E} \int_{\mathbb{R}} x^k d\mu_n &= \frac{1}{2^n} n^{k/2} 2^n \mathbb{E}[\mathbb{P}[X_v(k) = v]] \\ &= n^{k/2} \mathbb{E}[\mathbb{P}[X_v(k) = v]]\end{aligned}$$

In the following, C is a constant depending on k whose value may change from instance to instance. Let $m_0 = 8(k+1)\log_2 n$. By Lemma 1, with probability greater than $1 - Cn^{-k}$, the event F_v holds, that is, every vertex in $B(v,k)$ is contained in no more than Cm_0^{k+1} cycles of length at most k. By conditioning on F_n, we have

$$\begin{aligned}\mathbb{E}[\mathbb{P}[X_v(k) = v]] = {} & \mathbb{E}[\mathbb{P}[X_v(k) = v] | F_n] \mathbb{P}[F_n] \\ & + \mathbb{E}\left[\mathbb{P}[X_v(k) = v] | F_n^c\right] \mathbb{P}\left[F_n^c\right]\end{aligned}$$

The second term on the right-hand side is bounded below by 0 and above by

$$\mathbb{E}\left[\mathbb{P}[X_v(k) = v] | F_n^c\right] \mathbb{P}\left[F_n^c\right] \leq \mathbb{P}\left[F_n^c\right] \leq Cn^{-k} = o\left(n^{-k/2}\right)$$

Since $\mathbb{P}[F_n] = 1 - o(1)$, we then have

$$\begin{aligned}\mathbb{E}[\mathbb{P}[X_v(k) = v]] = {} & (1 + o(1)) \\ & \mathbb{E}[\mathbb{P}[X_v(k) = v] | F_n] + o\left(n^{-k/2}\right)\end{aligned}$$

To bound this term, we will count the number of random walks that return to the origin.

A step $(X_v(t), X_v(t+1))$ is said to be a *forward step* if $d_{\boldsymbol{G}_n}(v, X_v(t)) < d_{\boldsymbol{G}_n}(v, X_v(t+1))$, and a *backward step* if $d_{\boldsymbol{G}_n}(v, X_v(t)) \geq d_{\boldsymbol{G}_n}(v, X_v(t+1))$. By analyzing the combinatorics of forward and backward steps, it was shown by McKay [17, Lemma 2.1] that in an n-regular graph where every ball $B(v,k)$ has no cycles at all,

$$\begin{aligned}&\#\{\text{Walks of length } k \text{ which return to the origin}\} \\ &\qquad = (1 + o(1)) n^{k/2} C_{k/2}\end{aligned} \tag{19}$$

We now show that under F_n, the number of walks in $\boldsymbol{G}_n$ is of the same magnitude. Let ℓ be the number of forward steps of the walk $X_v(t)$ which are part of a cycle of length no larger than k, and suppose that $X_v(k) = v$.

If $\ell = 0$, then the walk must make $k/2$ forward steps and $k/2$ backward steps, since it returns to the origin. This means that k must be even, and the walk traces out a rooted tree with $k/2$ edges. Since the number of cycles with at most k edges is no larger than $C(\log n)^{k+1}$, there are at least $n - C(\log n)^{k+1} - 1$ choices for every forward step. By (19), the total number of walks with $\ell = 0$ is then equal to

$$(1+o(1))n^{k/2}C_{k/2}$$

when k is even, and 0 when k is odd.

If $\ell > 0$, then the walk makes ℓ forward steps which are part of a cycle, and no more than $k/2-\ell$ forward steps which are not part of a cycle. There are no more than k backward steps, and each such step has no more than $\left(C(\log n)^{k+1}+1\right)$ options. In total, the number of such walks is then bounded above by

$$(1+o(1))n^{k/2-\ell}\left(C(\log n)^{k+1}\right)^{k+\ell} = O\left(n^{k/2-\ell}(\log n)^{4k}\right)$$

Altogether, since the total number of walks of length k is n^k, we have

$$\begin{aligned} n^{k/2}&\mathbb{E}[\mathbb{P}[X_v(k)=v]] \\ &= n^{k/2}(1+o(1))\frac{1}{n^k}\left(n^{k/2}C_{k/2}+\sum_{l=1}^{k/2}O\left(n^{k/2-\ell}(\log n)^{4k^2}\right)\right) \\ &= (1+o(1))C_{k/2} \end{aligned}$$

as needed.

All that is left is to show that the variance is small. By (18), the second moment of $\int x^k d\mu_n$ is given by

$$\begin{aligned} \mathbb{E}&\left[\left(\int_{\mathbb{R}} x^k d\mu_n\right)^2\right] \\ &= \mathbb{E}\left[\left(\frac{1}{2^n}n^{k/2}\sum_{i=1}^{2^n}\mathbb{P}[X_i(k)=i]\right)^2\right] \\ &= \mathbb{E}\left[\frac{1}{2^{2n}}n^k\sum_{i,j=1}^{2^n}\mathbb{P}[X_i(k)=i]\mathbb{P}\left[X_j(k)=j\right]\right] \end{aligned}$$

Set $m_0 = 16(k+1)\log_2 n$. Recall that for a vertex $v \in V_n$, F_v is the event that each vertex in $B(v,k)$ is contained in no more than $C(m_0+1)^{k+1}$ cycles of length no more than k. Denote $F_{i,j} = F_i \cap F_j$. By Lemma 1, $\mathbb{P}\left[F_{i,j}\right] \ge 1 - 2C2^{-m_0}m_0^{2k+2}n^{2k+1}$. By the law of total probability, we have

$$\begin{aligned} \mathbb{E}&\left[\left(\int_{\mathbb{R}} x^k d\mu_n\right)^2\right] \\ &= \frac{1}{2^{2n}}\sum_{i,j=1}^{2^n}\mathbb{E}\left[n^k\mathbb{P}[X_i(k)=i]\mathbb{P}\left[X_j(k)=j\right]|F_{i,j}\right]\mathbb{P}\left[F_{i,j}\right] \\ &\quad + \frac{1}{2^{2n}}\sum_{i,j=1}^{2^n}\mathbb{E}\left[n^k\mathbb{P}[X_i(k)=i]\mathbb{P}\left[X_j(k)=j\right]|F_{i,j}^c\right]\mathbb{P}\left[F_{i,j}^c\right] \end{aligned}$$

The second term on the right-hand-side is bounded below by 0 and above, due to the choice of m_0, by $o(1)$. Thus

$$\begin{aligned} \mathbb{E}\left[\left(\int_{\mathbb{R}} x^k d\mu_n\right)^2\right] &= (1+o(1))\frac{1}{2^{2n}} \\ &\times \sum_{i,j=1}^{2^n}\mathbb{E}\left[n^k\mathbb{P}[X_i(k)=i]\mathbb{P}\left[X_j(k)=j\right]|F_{i,j}\right] + o(1) \end{aligned}$$

Using the same path-counting argument as above, by (19) we have that under $F_{i,j}$,

$$\mathbb{E}n^k\mathbb{P}[X_i(k)=i]\mathbb{P}\left[X_j(k)=j\right] = (1+o(1))C_{k/2}^2$$

and taking the sum overall i and j shows that

$$\mathbb{E}\left[\left(\int_{\mathbb{R}} x^k d\mu_n\right)^2\right] = (1+o(1))\left(\mathbb{E}\int_{\mathbb{R}} x^k d\mu_n\right)^2$$

which implies that $\mathrm{Var}\left(\int_{\mathbb{R}} x^k d\mu_n\right) \to 0$.

4.5 | Asymmetry

The proof of Theorem 5 relies on the following lemma, whose proof we postpone to the end of this section.

Lemma 3. *There exists a universal constant $C > 0$ such that the probability that there exists a decomposition of V_n into two disjoint subsets other than $V_n^0 \sqcup V_n^1$ such that the edges between them form a matching is smaller than $Cn2^{-n}$.*

Proof of Theorem 5. Let X_n be the number of automorphisms of $\boldsymbol{G}_n$. We partition these permutations into three kinds:

1. Automorphisms of $\boldsymbol{G}_n$ that swap between V_n^0 and V_n^1. Let W_n be the number of these automorphisms.
2. Automorphisms of $\boldsymbol{G}_n$ that preserve both V_n^0 and V_n^1. Let Y_n be the number of these automorphisms. Note that $Y_n \ge 1$, since it always counts the trivial automorphism.
3. Automorphisms of $\boldsymbol{G}_n$ that replace a proper subset $A_0 \subseteq V_n^0$ with a proper subset $A_1 \subseteq V_n^1$ of the same size (so that $V_n^0 \setminus A_0$ and $V_n^1 \setminus A_1$ stay inside V_n^0 and V_n^1, respectively). Let Z_n be the number of these automorphisms.

If φ is a non-trivial automorphism of the third kind, then the edges between A_0 and $V_n^0 \setminus A_0$ form a matching, and the edges between A_1 and $V_n^1 \setminus A_1$ form a matching (since, e.g., if there is a vertex $v \in A_0$ connected by more than one edge to $V_n^0 \setminus A_0$, then $\varphi(v)$ will have more than one edge across the main cut). But then, letting $A := A_0 \sqcup A_1$ and $B = V_n \setminus A$, we get that the edges between A and B form a matching as well, giving a partition $V_n = A \sqcup B$ with a matching between them. By Lemma 3, the probability that such a matching exists (and therefore, that there is a non-trivial automorphism swapping A_0 and A_1) is bounded by $O(n2^{-n})$. Thus, denoting by F the event $F := \{\exists m \in [n/20, n-1] \text{ s.t. } Z_m > 0\}$, we have

$$\mathbb{P}[F] = O\left(n^2 2^{-n/20}\right) \tag{20}$$

We turn to bound Y_n, W_n. For brevity, we abbreviate $\sigma := \sigma_{n-1}$. In the first two types, the values of an automorphism φ on V_n^0 determines the value of φ on all V_n. Explicitly, in the first case, for every $v \in V_{n-1}$, if we denote $\varphi(v,0) = (\varphi_0(v),1)$ and $\varphi(v,1) = (\varphi_1(v),0)$, then we must have $\varphi_0(v) = \sigma\varphi_1\sigma(v)$. For automorphisms of the second kind, we have similarly $\varphi_0 = \sigma^{-1}\varphi_1\sigma$. In both cases it must be that $\varphi_0, \varphi_1 \in \mathrm{Aut}(\boldsymbol{G}_{n-1})$. So in particular, $W_n, Y_n \le X_{n-1}$, and

$$W_n + Y_n \le 2X_{n-1} \tag{21}$$

We first bound $\mathbb{P}[W_n \ge 1 | X_{n-1}]$. By Markov's inequality, this is at most $\mathbb{E}[W_n | X_{n-1}]$. Write out $W_n = \sum_{\varphi_0} \sum_{\varphi_1} \mathbf{1}_{\varphi_0 \in \mathrm{Aut}(\boldsymbol{G}_{n-1})} \cdot \mathbf{1}_{\varphi_1 \in \mathrm{Aut}(\boldsymbol{G}_{n-1})} \cdot \mathbf{1}_{\varphi_0 = \sigma\varphi_1\sigma}$, and in particular we have

$$\begin{aligned} &\mathbb{E}[W_n | X_{n-1}] \\ &\quad \le X_{n-1}^2 \max_{\varphi_0, \varphi_1} \{\mathbb{P}_\sigma[\varphi_0 = \sigma\varphi_1\sigma] | \varphi_0, \varphi_1 \text{ bijections of } V_{n-1}\} \end{aligned}$$

So we need to bound $\mathbb{P}_\sigma[\varphi_0 = \sigma\varphi_1\sigma]$. Denote $A_0 = V_{n-1}$. For $v_0 \in A_0$, let E_{v_0} be the event that $\varphi_0(v_0) = \sigma\varphi_1\sigma(v_0)$. In order for E_{v_0} to hold, we must have either i) $E_{v_0} \cap \{\varphi_1\sigma(v_0) = v_0\}$ holds, or ii) $E_{v_0} \setminus \{\varphi_1\sigma(v_0) = v_0\}$ holds. The probability of the first event is at most the probability of $\varphi_1\sigma(v_0) = v_0$, which equals $1/2^{n-1}$, while the probability of the second given that $\varphi_1\sigma(v_0) \ne v_0$ is $\frac{1}{2^{n-1}-1}$. In particular E_{v_0} holds with probability no greater than $2/\left(2^{n-1}-1\right)$. Conditioned on E_{v_0}, the permutation σ is a uniform permutation over the set $A_1 = A_0 \setminus \{v_0, \varphi_1\sigma(v_0)\}$, with $|A_1| \ge |A_0| - 2$. By iteratively conditioning on $E_{v_0}, E_{v_1}, \ldots$, where $v_i \in A_i$, $i = 0, \ldots, 2^{n/3}-1$, we have that $\mathbb{P}[\varphi_0 = \sigma\varphi_1\sigma] \le \left(\frac{2}{2^{n-1}-2^{1+(n/3)}-1}\right)^{2^{n/3}} \le 2^{-(n-7)2^{n/3}}$. Hence

$$\mathbb{E}[W_n | X_{n-1}] \le 2^{-(n-7)2^{n/3}} \cdot X_{n-1}^2 \tag{22}$$

Next we bound $\mathbb{P}[Y_n > 1 | X_{n-1}, F^c]$. Although the equation $\varphi_0 = \sigma^{-1}\varphi_1\sigma$ seems similar to the analogous equation $\varphi_0 = \sigma\varphi_1\sigma$ for W_n, we shouldn't expect the same argument to hold, since (for example) even if $X_{n-1} = 1$, we expect $W_n = 0$, whereas $Y_n \ge 1$ always since it counts the identity. The problem lies with automorphisms with small conjugacy classes. For a given φ_1 and uniformly random σ, the element $\sigma^{-1}\varphi_1\sigma$ is a uniform element in the conjugacy class of φ_1. The probability $\mathbb{P}_\sigma\left[\varphi_0 = \sigma^{-1}\varphi_1\sigma\right]$ is then bounded by one over the size of the conjugacy class of φ_0 (it is 0 if φ_0 and φ_1 are not conjugate). The following claim, whose proof is found at the end of the section, shows that under F^c, these classes must be large.

Claim 2. Assume that F^c occurs. Then the conjugacy class for every $\mathrm{Id} \ne \varphi \in \mathrm{Aut}(\boldsymbol{G}_{n-1})$ has size at least $2^{\frac{1}{4}n2^{n/4}}$.

As in the case of W_n, we have

$$\begin{aligned} &\mathbb{P}[Y_n > 1 | X_{n-1}, F^c] \\ &\quad \le \mathbb{E}[Y_n - 1 | X_{n-1}, F^c] \\ &\quad \le \mathbb{E}\left[X_{n-1}^2 | F^c\right] \max_{\varphi_0, \varphi_1 \in \mathrm{Aut}(\boldsymbol{G}_{n-1}) \setminus \{\mathrm{Id}\}} \mathbb{P}_\sigma\left[\varphi_0 = \sigma^{-1}\varphi_1\sigma\right] \\ &\quad \le \mathbb{E}\left[X_{n-1}^2 | F^c\right] 2^{-\frac{1}{4}n2^{n/4}} \end{aligned}$$

where the last inequality is due to Claim 2, since the maximum is taken over elements with a conjugacy class of size at least $2^{\frac{1}{4}n2^{n/4}}$.

Finally, under F^c, the only possible automorphisms for $m \in [n/20, n-1]$ are of the first two kinds, and by (21) we have

$$X_{n-1} \le 2^{n-n/20-1} X_{n/20} \le 2^{19n/20}\left(2^{n/20}\right)! \le 2^{n+\frac{1}{20}n2^{n/20}} \tag{23}$$

Thus

$$\begin{aligned} &\mathbb{P}[X_n > 1] \\ &\le \mathbb{P}[F] + \mathbb{P}[F^c \cap \{X_n > 1\}] \\ &\le \mathbb{P}[F] + \mathbb{P}[Z_n > 0] + \mathbb{P}[F^c \cap \{W_n > 0\}] + \mathbb{P}[F^c \cap \{Y_n > 1\}] \\ &\le \mathbb{P}[F] + \mathbb{P}[Z_n > 0] + \mathbb{P}\left[W_n > 0, X_{n-1} \le 2^{n+\frac{1}{20}n2^{n/20}}, F^c\right] \\ &\quad + \mathbb{P}\left[Y_n > 1, X_{n-1} \le 2^{n+\frac{1}{20}n2^{n/20}}, F^c\right] \\ &\le \mathbb{P}[F] + \mathbb{P}[Z_n > 0] + \frac{2^{2n+\frac{1}{10}n2^{n/20}}}{2^{(n-7)2^{n/3}}} + \frac{2^{2n+\frac{1}{10}n2^{n/20}}}{2^{\frac{1}{4}n2^{n/4}}} = O\left(n^2 2^{-n/20}\right) \end{aligned}$$

as needed.

Proof of Lemma 3. We start with some preliminaries which will be of use later on in the proof. Let $V_n = \boldsymbol{A} \sqcup \boldsymbol{B}$ be a uniformly random partition of V_n into two halves of equal size, and let $V_n = \boldsymbol{A}' \sqcup \boldsymbol{B}'$ be a partition where $\boldsymbol{A}'$ is a binomial random subset of V_n with success probability $1/2$. The difference between these two random partitions can be quantified as follows: for any arbitrary set Σ of equal-sized partitions of V_n, we have

$$\begin{aligned} \mathbb{P}[(\boldsymbol{A}, \boldsymbol{B}) \in \Sigma] &= \mathbb{P}\left[\left(\boldsymbol{A}', \boldsymbol{B}'\right) \in \Sigma \,\middle|\, |\boldsymbol{A}'| = 2^{n-1}\right] \qquad (24) \\ &\le \frac{\mathbb{P}\left[\left(\boldsymbol{A}', \boldsymbol{B}'\right) \in \Sigma\right]}{\mathbb{P}\left[|\boldsymbol{A}'| = 2^{n-1}\right]} \qquad (24) \end{aligned}$$

The denominator in the right-hand side can be approximated by the de Moivre–Laplace limit theorem, which states that

$$\mathbb{P}\left[|\boldsymbol{A}'| = 2^{n-1}\right] = \frac{1+o(1)}{\sqrt{\pi}} 2^{(1-n)/2} \tag{25}$$

Note that V_n contains a vertex-disjoint union of 2^{n-2} copies of P_3, the 3 vertex path, and let Σ be the set of all equal-sized partitions which do not separate the middle vertex from the other two vertices of any of these paths. When the vertices are partitioned randomly and independently, the number of paths split this way is a binomial random variable with parameters 2^{n-2} and $1/4$, and thus $\mathbb{P}\left[\left(\boldsymbol{A}', \boldsymbol{B}'\right) \in \Sigma\right] \le e^{-\ln\frac{4}{3}\cdot 2^{n-2}}$, and so by (24) and (25), we have

$$\mathbb{P}[(\boldsymbol{A}, \boldsymbol{B}) \in \Sigma] \le e^{-\ln\frac{4}{3}\cdot 2^{n-2}} \frac{\sqrt{\pi}}{1+o(1)} 2^{(n-1)/2} \le c \cdot e^{n-\ln\frac{4}{3}\cdot 2^{n-2}} \tag{26}$$

for some constant $c > 0$. Let us now choose C so large that the lemma is true for all $n \le n_0$ for some large enough n_0. We proceed by induction on n. Assume that the lemma is true for $\boldsymbol{G}_{n-1}$. For a decomposition $V_n = A \sqcup B$ other than $V_n^0 \sqcup V_n^1$, let $E_n(A, B)$ be the event that the edges between A and B form a matching. Let $p_n := \mathbb{P}[\exists A, B \text{ s.t. } E_n(A, B)]$. We will show that $p_{n-1} \le C(n-1)2^{-(n-1)}$ implies $p_n \le Cn2^{-n}$.

Fix a decomposition $V_n = A \sqcup B$ and set $A_j = V_n^j \cap A$ and $B_j = V_n^j \cap B$ for $j = 0, 1$, so that $A_0 \sqcup B_0 = V_n^0$ and $A_1 \sqcup B_1 = V_n^1$. We consider three cases.

1. If both cuts coincide with the cuts of the $(n-1)$th generation (i.e., $A_0 \sqcup B_0 = V_n^{00} \sqcup V_n^{10}$ and $A_1 \sqcup B_1 = V_n^{01} \sqcup V_n^{11}$), then for $A \sqcup B$ to induce a matching, all vertices of A_0

should send the edges of the last generation to A_1, and all vertices of B_0 should send the edges of the last generation to B_1. The probability of this event is exactly $\binom{2^{n-1}}{2^{n-2}}^{-1}$.

2. If both cuts differ from the $(n-1)$th generation cuts, then assume first that A_0 is empty. By connectivity of $\boldsymbol{G}_{n-1}$, there exists an edge between some vertex $a \in A_1$ and a vertex $b \in B_1$. Since a also sends an edge to B_0 (induced by σ_{n-1}), in this case, the edges do not form a matching. We can therefore assume that all of A_0, B_0, A_1, B_1 are non-empty. Assume without loss of generality that $|B_1| \geq |A_1|$. Let $a \in A_0$ be a vertex that sends an edge to B_0 (again, such a vertex exists by connectivity of G_{n-1}). Since $|B_1| \geq |A_1|$, the probability that there are no edges from a to B_1 induced by σ_{n-1} is at most $1/2$, so with probability greater than $1/2$ we do not get a matching. We get that

$$\mathbb{P}\left[\exists A, B \text{ s.t. } \left\{A_0 \sqcup B_0 \neq V_n^{00} \sqcup V_n^{10}, A_1 \sqcup B_1 \neq V_n^{01} \sqcup V_n^{11}\right\} \cap E_n(A, B)\right] \leq \frac{1}{2} p_{n-1}$$

3. Finally, if, say, the cut $A_0 \sqcup B_0$ coincides with the respective $(n-1)$th generation cut $V_n^{00} \sqcup V_n^{10}$, and $A_1 \sqcup B_1 \neq V_n^{01} \sqcup V_n^{11}$, then $A_1 \sqcup B_1$ should divide the set V_1 into halves; otherwise (say, if $|A_1| > |B_1|$), B_0 sends at least one nth generation edge to A_1, and so there is a vertex in B_0 with at least two neighbors in A, and we do not get a matching. Moreover, we may also claim that the nth generation edges form a matching between A_0 and A_1, and between B_0 and B_1 (since there is a matching between A_0 and B_0, A_0 cannot have an edge with B_1, and B_0 cannot have an edge with A_1). Then the desired probability is exactly the probability that the cut $A_1 \sqcup B_1$ of $\boldsymbol{G}_{n-1}$ forms a matching. Since the ends of the edges of the matching between V_n^0 and V_n^1 with first vertices in sets V_n^{00} and V_n^{10} form a decomposition of $\boldsymbol{G}_{n-1}$ into halves which is independent of $\boldsymbol{G}_{n-1}$ itself, the sets A_1 and B_1 are a uniformly random partition of V_n^1. The probability that a uniformly random balanced cut of $\boldsymbol{G}_{n-1}$ is a matching is at most the probability that this cut does not separate the middle vertex of any P_3, which is bounded by $c \cdot 2^{n-1-2^{n-3}}$ due to (26).

Putting all these together, we get

$$p_n \leq \frac{1}{2} p_{n-1} + \binom{2^{n-1}}{2^{n-2}}^{-1} + c \cdot e^{n-1-\ln\frac{4}{3}\cdot 2^{n-3}} < Cn2^{-n}$$

for $n > n_0$ large enough.

Proof of Claim 2. Assume that F^c holds, that is, every automorphism of $\boldsymbol{G}_m$ either swaps or preserves V_m^0, V_m^1 for $m \in [n/20, n-1]$. We first show that every non-identity $\varphi \in \mathrm{Aut}(\boldsymbol{G}_{n-1})$ has at least $2^{19n/20}$ points that are not fixed.

Let $m \in [n/20, n-1]$, and let $\psi \in \mathrm{Aut}(\boldsymbol{G}_m)$. Since F^c holds, ψ either swaps or preserves V_m^0 and V_m^1, and so can be represented by the pair (ψ_0, ψ_1) as above. If it swaps V_m^0 and V_m^1, then it has no fixed points. Hence, if ψ has any fixed points, it must be preserving, and its fixed points are a union of the fixed points of ψ_0 and ψ_1. In this case, ψ_0 and ψ_1 are conjugate, so they have the same number of fixed points; in particular, the number of fixed points (resp. non-fixed points) of ψ is equal to twice the number of fixed points (resp. non-fixed points) of ψ_0.

Thus by induction, if $\varphi \in \mathrm{Aut}(\boldsymbol{G}_{n-1})$ has any fixed points, then the number of non-fixed points is equal to $2^{n-1-n/20}$ times the number of non-fixed points of any of its restrictions $\psi := \varphi_{|V_n^z}$, where $z \in \{0,1\}^{n-1-n/20}$. If ψ is the identity, then φ is the identity also. Otherwise, ψ has at least 2 non-fixed points, and so φ has at least $2^{19n/20}$ non-fixed points on $\boldsymbol{G}_{n-1}$.

Next we get our bound for the size of the conjugacy class of φ. Recall that we can express φ as a composition of disjoint cycles.

1. If φ has a cycle of length $m \geq 2^{2n/5} + 1$, then the number of conjugacy classes is bounded below by the number of conjugacy classes where (say) 1 is in such a cycle. The number of such permutations is at least $\binom{2^{n-1}-1}{m-1} \cdot (m-1)!$ (since we need to choose $m-1$ more elements, and then order them in a cycle together with 1). This is $(2^{n-1}-1) \cdot (2^{n-1}-2) \cdots (2^{n-1}-m) \geq 2^{(n-2)2^{2n/5}}$.
2. Otherwise, the maximal cycle length is at most $2^{2n/5}$. We have at least $2^{19n/20}$ points which are not fixed, so there are at least $r = 2^{11n/20}$ cycles. The number of conjugacy classes is then lower-bounded by the number of conjugacy classes where all the elements $1, 2, \ldots, r$ are in distinct cycles, and $r+1, r+2, \ldots, 2r$ are in the same distinct cycles. For every fixed choice of cycles for the first r elements, there are $r!$ ways to choose where to put $r+1, r+2, \ldots, 2r$. This is at least $2^{11n/20}! \geq 2^{\frac{1}{4}n2^{n/4}}$. □

Acknowledgments

We thank Elad Tzalik for discussions about the diameter and open questions, and Sahar Diskin for finding some of the references to existing literature. I.B. is supported by the Israel Science Foundation. R.G. is supported by the Adams Fellowship Program of the Israel Academy of Sciences and Humanities, the European Research Council and by the Israel Science Foundation. M.Z. is supported in part by the Israel Science Foundation Grant 2110/22.

Data Availability Statement

Data sharing is not applicable to this article as no new data were created or analyzed in this study.

Endnotes

[1] In Question 2 of Section 3, we ask whether the bound in Theorem 2 can be improved to $\frac{n}{\log_2 n}(1+o(1))$ for general random twisted hypercubes. After the submission of this manuscript, Aragão et al. [15] indeed proved that $D(\boldsymbol{G}_n) = (1+o(1))\frac{n}{\log_2 n}$ with high probability, regardless of the joint distribution of the copies.

[2] After the submission of our manuscript, this question was solved in the affirmative by Aragão et al. [15].

References

1. S. Abraham and K. Padmanabhan, "The Twisted Cube Topology for Multiprocessors: A Study in Network Asymmetry," *Journal of Parallel and Distributed Computing* 13, no. 1 (1991): 104–110.

2. F. B. Chedid and R. B. Chedid, "A New Variation on Hypercubes With Smaller Diameter," *Information Processing Letters* 46, no. 6 (1993): 275–280.

3. P. A. J. Hilbers, M. R. J. Koopman, and J. L. A. van de Snepscheut, "The Twisted Cube," in *PARLE Parallel Architectures and Languages Europe*, eds. J. W. de Bakker, A. J. Nijman, and P. C. Treleaven (Berlin, Germany: Springer, 1987), 152–159.

4. A. Dudek, X. Pérez-Giménez, P. Prałat, H. Qi, D. West, and X. Zhu, "Randomly Twisted Hypercubes," *European Journal of Combinatorics* 70 (2018): 364–373.

5. D. A. Levin and Y. Peres, *Markov Chains and Mixing Times* (Providence, RI: American Mathematical Society, 2017).

6. F. Harary, J. P. Hayes, and H.-J. Wu, "A Survey of the Theory of Hypercube Graphs," *Computers & Mathematics with Applications* 15, no. 4 (1988): 277–289.

7. S. G. Bobkov, "An Isoperimetric Inequality on the Discrete Cube, and an Elementary Proof of the Isoperimetric Inequality in Gauss Space," *Annals of Probability* 25, no. 1 (1997): 206–214.

8. Z. Jiang and A. Yehudayoff, "An Isoperimetric Inequality for Hamming Balls and Local Expansion in Hypercubes," *Electronic Journal of Combinatorics* 29, no. 1 (2022): P1.15.

9. R. O'Donnell, *Analysis of Boolean Functions* (New York, NY: Cambridge University Press, 2014).

10. B. Bollobás and W. Fernandez de la Vega, "The Diameter of Random Regular Graphs," *Combinatorica* 2, no. 2 (1982): 125–134.

11. R. Bauerschmidt, J. Huang, A. Knowles, and H.-T. Yau, "Edge Rigidity and Universality of Random Regular Graphs of Intermediate Degree," *Geometric and Functional Analysis* 30, no. 3 (2020): 693–769.

12. E. Lubetzky and A. Sly, "Cutoff Phenomena for Random Walks on Random Regular Graphs," *Duke Mathematical Journal* 153, no. 3 (2010): 475–510.

13. I. Dumitriu and S. Pal, "Sparse Regular Random Graphs: Spectral Density and Eigenvectors," *Annals of Probability* 40, no. 5 (2012): 2197–2235.

14. P. Gao, "The Number of Perfect Matchings, and the Nesting Properties, of Random Regular Graphs," arXiv preprint arXiv:2104.11850 (2021).

15. L. Aragão, M. Collares, G. Dahia, and J. P. Marciano, "The Diameter of Randomly Twisted Hypercubes," arXiv preprint arXiv:2306.01728 (2021).

16. A. E. Brouwer and W. H. Haemers, *Spectra of Graphs* (New York, NY: Springer, 2012).

17. B. D. McKay, "The Expected Eigenvalue Distribution of a Large Regular Graph," *Linear Algebra and Its Applications* 40 (1981): 203–216.

18. B. D. McKay and N. C. Wormald, "Automorphisms of Random Graphs With Specified Vertices," *Combinatorica* 4, no. 4 (1984): 325–338.

19. A. W. Marshall, *Inequalities: Theory of Majorization and Its Applications*, 2nd ed. (New York, NY: Springer, 2011).

20. X. Zhu, "A Hypercube Variant With Small Diameter," *Journal of Graph Theory* 85, no. 3 (2017): 651–660.

21. D. Ellis, "Almost Isoperimetric Subsets of the Discrete Cube," *Combinatorics, Probability and Computing* 20, no. 3 (2011): 363–380.

22. A. Louis, P. Raghavendra, and S. Vempala, "The Complexity of Approximating Vertex Expansion," in *2013 IEEE 54th Annual Symposium on Foundations of Computer Science* (New York, NY: IEEE, 2013), 360–369.

23. F. P. Preparata and J. Vuillemin, "The Cube-Connected Cycles: A Versatile Network for Parallel Computation," *Communications of the ACM* 24, no. 5 (1981): 300–309.

24. I. Benjamini and O. Schramm, "Recurrence of Distributional Limits of Finite Planar Graphs," *Electronic Journal of Probability* 6, no. 23 (2001): 1–13.

25. J. H. van Lint and R. M. Wilson, *A Course in Combinatorics*, 2nd ed. (Cambridge, UK: Cambridge University Press, 2001).

26. G. Kemkes, X. Pérez-Giménez, and N. Wormald, "On the Chromatic Number of Random d-Regular Graphs," *Advances in Mathematics* 223, no. 1 (2010): 300–328.