Electronic Colloquium on Computational Complexity, Report No. 89 (2024)

Tight Bounds for the Zig-Zag Product

Gil Cohen* Itay Cohen† Gal Maor‡

May 8, 2024

Abstract

The Zig-Zag product of two graphs, $Z = G \textcircled{z} H$, was introduced in the seminal work of Reingold, Vadhan, and Wigderson (Ann. of Math. 2002) and has since become a pivotal tool in theoretical computer science. The classical bound, which is used throughout, states that the spectral expansion of the Zig-Zag product can be bounded roughly by the sum of the spectral expansions of the individual graphs, $\omega_Z \leq \omega_H + \omega_G$.

In this work we derive, for every (vertex-transitive) c-regular graph H on d vertices, a *tight* bound for ω_Z by taking into account the *entire* spectrum of H. Our work reveals that the bound, which holds for every graph G, is precisely the minimum value of the function

$$\frac{x}{c^2} \cdot \sqrt{1 - \frac{d \cdot h(x)}{x \cdot h'(x)}}$$

in the domain (c^2, ∞), where $h(x)$ is the characteristic polynomial of H^2. As a consequence, we establish that Zig-Zag products are indeed intrinsically quadratic away from being Ramanujan.

We further prove tight bounds for the spectral expansion of the more fundamental replacement product. Our lower bounds are based on results from analytic combinatorics, and we make use of finite free probability to prove their tightness. In a broader context, our work uncovers intriguing links between the two fields and these well-studied graph operators.

*Tel Aviv University. `gil@tauex.tau.ac.il`. Supported by ERC starting grant 949499 and by the Israel Science Foundation grant 1569/18.

†Tel Aviv University. `itay3@mail.tau.ac.il`. Supported by ERC starting grant 949499.

‡Tel Aviv University. `galmaor@mail.tau.ac.il`. Supported by ERC starting grant 949499.

ISSN 1433-8092

Contents

1 Introduction

Expander graphs are fundamental across a wide range of fields within theoretical computer science as well as in coding theory and cryptography, among others. Their significance largely arises from the flexibility to view expansion through different lenses, whether combinatorial, probabilistic, or linear algebraic. This multifaceted understanding offers a unique advantage, allowing for the deduction of combinatorial properties of graphs by examining the spectral properties of related operators such as their adjacency matrices. The prominent example of this interplay is seen in the study of the *spectral expansion* whose definition we briefly recall next.

Let G be an undirected d-regular graph on n vertices, referred to as an (n, d)-graph throughout, and let $\mathbf{A}$ be its adjacency matrix. Since G is undirected, $\mathbf{A}$ is symmetric and so its spectrum is real-valued. We denote the eigenvalues of $\mathbf{A}$ by $d = \lambda_1 \geq \lambda_2 \geq \cdots \geq \lambda_n$. The *spectral expansion* of G, denoted as $\lambda(G)$, is given by $\max(\lambda_2, |\lambda_n|)$. We further denote the normalized spectral expansion of G by $\omega(G) = \frac{\lambda(G)}{d} \in [0, 1]$. We alternate between the two variants–the normalized and the unnormalized–depending on context.

An *expander* is a graph G which is, informally speaking, a sparse approximation of the complete graph. Formally, one requires that the normalized spectral expansion $\omega(G)$ is bounded away from 1 [1]. For a typical application, one "pays" a cost that increases with the degree d and has an "error" that vanishes as $\omega(G)$ tends to 0. This raises the question of what is the best tradeoff between $\omega(G)$ and the degree d of G. From the Alon-Boppana bound [Nil91], which is usually stated in terms of $\lambda(G)$, it follows that for every $\varepsilon > 0$ there are only finitely many d-regular graphs G with $\lambda(G) \leq 2\sqrt{d-1} - \varepsilon$. A d-regular graph G satisfying $\lambda(G) \leq 2\sqrt{d-1}$ is called a *Ramanujan graph*.

1.1 The Zig-Zag product

In their landmark paper [RVW00], Reingold, Vadhan, and Wigderson introduced the Zig-Zag product, a novel approach that provides a simple combinatorial method for constructing expander graphs. The combinatorial nature of the Zig-Zag product and its versatility rendered it exceedingly valuable. Indeed, soon after the publication of [RVW00], Reingold [Rei08] obtained his seminal result, $\mathbf{SL} = \mathbf{L}$, based on the Zig-Zag product. In this case, the Zig-Zag product was not used to construct expander graphs per se but rather to "transform" an existing graph into an expander, crucially without altering the structure of its connected components.

Since then, the Zig-Zag product has gained significant attention. Many works have utilized the Zig-Zag product, while others have adopted the underlying theme, which Reingold has dubbed the "Zig-Zag recipe" (see Goldreich's insightful survey [Gol05]). Prominent ex-

[1]It is common in this context to consider an infinite family of graphs. However, we will exclude this technical detail from our informal discussion for simplicity.

amples include Dinur's proof of the PCP theorem [Din07], the state-of-the-art construction of locally decodable and correctable codes [KMRZS17], and the construction of near-optimal small-bias sets by Ta-Shma [TS17] which in turn is based on the improved Zig-Zag variant dubbed the *wide-replacement product* due to Ben-Aroya and Ta-Shma [BATS11].

We turn to describe the Zig-Zag product and its known properties. Our discussion here will be somewhat informal, with a more formal treatment available in Section 4.1. Let $G = (V_G, E_G)$ be an (n, d)-graph. We assume that every vertex $v \in V_G$ labels its neighbors using the label set $[d] \triangleq \{1, \ldots, d\}$ in such a way that every vertex uses all d labels, though there is no requirement for label consistency across different vertices. Let $H = (V_H, E_H)$ be a (d, c)-graph, where we identify V_H with $[d]$. The Zig-Zag product of G and H, denoted as $G \textcircled{z} H$, is the graph on the vertex set $V_G \times V_H$, defined informally as follows: for every vertex $v \in V_G$, the d vertices $\{v\} \times [d]$ in the new graph are dubbed the *cloud* of v. Every edge adjacent to the vertex (v, i) in $G \textcircled{z} H$ is induced by the following walk: starting at (v, i) we take one step ("zig") of the form $(v, i) \to (v, \ell_1)$, where ℓ_1 is a neighbor of i in H, remaining within the same cloud; we follow by the unique step $(v, \ell_1) \to (u, \ell_2)$ where u is the ℓ_1 neighbor of v and v is the ℓ_2 neighbor of u; lastly, we take an additional H step ("zag") of the form $(u, \ell_2) \to (u, j)$ in u-s cloud, j being a neighbor of ℓ_2 in H. Note that $G \textcircled{z} H$ is an undirected graph on nd vertices of degree c^2.

The key result of [RVW00] is a bound on the spectral expansion of the Zig-Zag product. It is shown that for G and H which are ω_G and ω_H expanders, respectively, it holds that $G \textcircled{z} H$ has spectral expansion [2]

$$\omega(G \textcircled{z} H) \leq \omega_G + \omega_H + \omega_H^2. \tag{1}$$

Given the importance of the Zig-Zag product in theoretical computer science, it is pivotal to gain a better understanding on the spectral expansion of $G \textcircled{z} H$. Of particular interest is the setting in which H is taken to be a Ramanujan graph, as in many applications one may choose H, thus optimizing the spectral expansion of H is natural. In such case, $\omega_H = \Theta(\frac{1}{\sqrt{c}})$. Recall that the degree of $G \textcircled{z} H$ is $D = c^2$. Thus, according to the bound from Equation (1), all that can be said about the resulted graph is that it has spectral expansion $O(\frac{1}{D^{1/4}})$ – quadratic away from Ramanujan.

1.2 Derandomized squaring

The question of understanding the spectral expansion associated with important graph operators has been asked, and partially answered, for a graph product which is related to the Zig-Zag product - the derandomized squaring. We turn to briefly discuss this operator which was introduced by Rozenman and Vadhan [RV05] for their re-derivation of Reingold's

[2] As a matter of fact, [RVW00] obtained a slightly stronger bound on the spectral expansion of the resulted graph (see Section 5.5.2) though this does not affect our discussion.

theorem. Informally, given an (n,d)-graph G and a (d,c)-graph H, the *derandomized square*, denoted $G \textcircled{s} H$, is obtained by placing a copy of H on the neighborhood of each vertex of G. The logic being that G^2 is obtained by placing the complete graph on every such neighborhood. As H, being an expander, serves as a sparse approximation of the complete graph, it is plausible that the resulted graph, $G \textcircled{s} H$, is a sparse approximation of G^2. In particular, its spectral expansion should approximate ω_G^2, where the quality of the approximation improves as $\omega_H \to 0$. Formally, Rozenman and Vadhan established the bound

$$\omega(G \textcircled{s} H) \leq \omega_G^2 + \omega_H, \tag{2}$$

where, note, the degree of $G \textcircled{s} H$ is cd (compared to a degree d^2 for G^2 and degree c^2 of $G \textcircled{z} H$). Further, the authors proved that their bound is tight in the sense that it cannot be improved if the only information that is incorporated to the bound is the spectral expansion of the two graphs.

The question of what can be said about the spectral expansion of $G \textcircled{s} H$ was addressed in [CCMP23], where the authors proved a lower bound on $\omega(G \textcircled{s} H)$, depending on the *entire* spectrum of H. Surprisingly, [CCMP23] also gave evidence, though not a proof, for the fact that the derandomized squaring operation is essentially as strong as possible, namely, $G \textcircled{s} H$ can get arbitrarily close to being Ramanujan (in particular, $\omega(G \textcircled{s} H) = O\left(\frac{1}{\sqrt{cd}}\right)$ is achievable). Thus, we have strong reasons to believe that the derandomized squaring operation is, in a sense, optimal, though the bound obtained by utilizing only the spectral expansions as given in Equation (2) is far from capturing this as it yields a bound no better than $\omega_H \approx \frac{1}{\sqrt{c}}$. Does the same hold true for the Zig-Zag product?

2 Our results

In this paper, we obtain a deeper insight into the Zig-Zag product, providing a definite answer to the previously mentioned question by examining the *entire* spectrum of H. We further derive tight bounds for another foundational graph-theoretic operation–the *replacement product*. This discussion will follow shortly. Our study of both the Zig-Zag and replacement products, along with the previously analyzed derandomized squaring operation [CCMP23], unveils fascinating interrelations among these combinatorial operators, finite free probability, and analytic combinatorics. We initially present our results in a relatively informal manner, postponing the detailed formal statements to later sections, specifically Sections 2.1 to 2.3.

Lower bound on the spectral expansion of the Zig-Zag product. Our first result focuses on the limitations of the Zig-Zag product, with the assumption throughout this paper that H is vertex-transitive (interestingly, we will not rely on this assumption in our upper bounds). Specifically, we establish a lower bound for the spectral expansion of $G \textcircled{z} H$. This bound considers the entire spectrum of the (d,c)-graph H, captured by the characteristic polynomial of its square, $\chi_x(H^2)$, and is applicable to *every* graph G. Our work reveals that the following function emerges as critically important in the analysis of the Zig-Zag product:

$$\mathcal{Z}_H(x) = x \cdot \sqrt{1 - \frac{d \cdot \chi_x(H^2)}{x \cdot \chi'_x(H^2)}}. \tag{3}$$

Indeed, our lower bound is given by

$$\lambda(G \textcircled{z} H) \geq \min_{x > c^2} \mathcal{Z}_H(x) - o_n(1),$$

where n is the number of vertices in G. The presence of H^2, rather than H, is perhaps expected given that defining an edge in the Zig-Zag product requires taking two steps on H. Our proof techniques lean on results from analytic combinatorics and the symbolic method.

This result reduces the difficult combinatorial problem of lower bounding the spectral expansion of Zig-Zag-ing with H to a straightforward minimization problem. Taking for example $H = C_6$, the length-6 cycle, we have that $\mathcal{Z}_{C_6}(x) = \sqrt{\frac{2x(x-2)}{x-3}}$, whose minimum in $(4, \infty)$ can be easily derived, $\sqrt{6} + \sqrt{2} \approx 3.86$. This should be compared to the Alon-Boppana bound for degree-4 graphs, $2\sqrt{3} \approx 3.46$. Naturally, for more complicated graphs, the minimization problem becomes more difficult though certainly approachable, and far easier than tackling the problem combinatorially. For example, the function corresponding to the Petersen graph is given by $\mathcal{Z}_{\mathrm{Pet}}(x) = \sqrt{\frac{3x(x^2-9x+12)}{x^2-11x+22}}$, from which we can deduce that $\lambda(G \textcircled{z} \mathrm{Pet}) \geq 7.11$. (see Section 5.5.2 for the exact bound).

Zig-Zag is inherently quadratically far from Ramanujan. We utilize the above result to deduce a universal lower bound for the spectral expansion of the Zig-Zag product with *any* c-regular vertex-transitive graph H. We prove that the classical bound obtained by [RVW00] is asymptotically optimal, meaning that $\omega(G \textcircled{z} H) = \Omega(\frac{1}{D^{1/4}})$ where, recall, $D = c^2$. Thus, the quadratic gap from Ramanujan is an intrinsic characteristic of the Zig-Zag operation. This is a striking contrast to the strong evidence gathered for the closely related derandomized squaring operation [CCMP23].

Proof of the tightness of the lower bound. Our lower bound acts as an analogue of the Alon-Boppana bound for Zig-Zag products. We complete the analysis by proving the existence of graphs of every size meeting it on the positive side of the spectrum, effectively acting as the Zig-Zag analogue of (one-sided) Ramanujan graphs. More formally, we prove that for every (d,c)-graph H–not necessarily vertex-transitive–and every integer n, there exists an (n,d)-graph G such that

$$\lambda_2(G \textcircled{z} H) \leq \min_{x>c^2} \mathcal{Z}_H(x). \tag{4}$$

Our proof utilizes *finite free probability*, following the methodology of the seminal works of Marcus, Spielman and Srivastava [MSS22, MSS18], who used it to prove the existence of bipartite Ramanujan graphs. A key observation in our proof is the need to diverge from the aforementioned works and consider the graph G in the *configuration model*. We stress that for derandomized squaring, proving a tightness result remains open [CCMP23].

Tight bounds for the replacement product. Another graph operation which in fact predates the Zig-Zag product is the *replacement product* denoted $G \textcircled{r} H$. Informally, one can think of the replacement product in the following way: every vertex v of G is replaced by a copy of H which we dub the *cloud* of v, where the original edges from G are connecting the clouds according to the edges of G (see Definition 4.3). This operator is more straightforward from the combinatorial perspective and often serves as an introductory concept for grasping the Zig-Zag product. Interestingly, [RVW00] used their bound on $\omega(G \textcircled{z} H)$ to prove a bound on $\omega(G \textcircled{r} H)$. We employ a similar approach to our work on the Zig-Zag product, proving a lower bound on the spectral expansion of $G \textcircled{r} H$ using analytic combinatorics, and prove its tightness using finite free probability. While the free analysis bears a striking resemblance to the Zig-Zag case (interestingly, informally speaking, the replacement product turns out to be an additive counterpart to the multiplicative Zig-Zag), the symbolic and analytic techniques we apply for the lower bound prove to be significantly more intricate.

2.1 Lower bound for the Zig-Zag product

In this section, we formally present our lower bound for the spectral expansion of Zig-Zag products. As in previous discussions, H is a vertex-transitive (d,c)-graph. There are certain degenerate choices for H which lead to trivial outcomes when used in Zig-Zag products, and are excluded from our results. These can be categorized into two distinct types: (1) graphs consisting solely of self-loops and a perfect matching, possibly with parallel edges; and (2) graphs composed of the disjoint union of 4-cycles (see Definition 5.1 and the short discussion following it). Any other graph H is deemed *Zig-Zag good*, or simply *good*.

Theorem 2.1. *Let H be a good vertex-transitive (d,c)-graph with $d \geq 3$ and $c \geq 2$. Then, for every (n,d)-graph G,*

$$\lambda(G \textcircled{z} H) \geq \min_{x>c^2} \mathcal{Z}_H(x) - o_n(1).$$

The full statement of the above theorem is given in Section 5.1. We give an example of its usage with H being the Petersen graph in Section 5.5.2, and provide a proof sketch of Theorem 2.1 in Section 3.2.

2.1.1 The universal Zig-Zag lower bound

At first glance, it appears that Theorem 2.1 cannot be applied to establish a universal bound on the spectral expansion of Zig-Zag products involving all vertex-transitive graphs of a given degree c, as it necessitates solving a distinct minimization problem for each graph H without offering a general bound akin to Equation (1). To address this challenge, we substitute the graph H with its universal cover, the infinite c-ary tree, denoted as $\mathcal{T}_c$. This approach is based on the premise that a graph's spectrum encodes the number of cycles of any specified length, making it reasonable to examine a graph with fewer cycles to derive a lower bound on spectral expansion. However, the infinite nature of $\mathcal{T}_c$ means it lacks an associated characteristic polynomial, rendering the direct application of Theorem 2.1 moot. To circumvent this issue, we restate the theorem in terms of the Cauchy transform,

$$\mathcal{G}_H(x) = \frac{1}{d} \cdot \frac{\chi_x'(H)}{\chi_x(H)} = \frac{1}{d} \sum_{i=1}^{d} \frac{1}{x - \lambda_i}, \tag{5}$$

and utilize the (continuous) Cauchy transform corresponding to $\mathcal{T}_c$, which takes the form $\mathcal{G}_{\mathcal{T}_c}(x) = \int \frac{1}{x-\lambda} \mu_c(\lambda)$ for some suitable measure μ_c known as the Kesten-McKay distribution with parameter c. Through this method, we establish the following theorem.

Theorem 2.2. *Let $d, c \geq 3$. For every (n, d)-graph G and every good vertex-transitive (d, c)-graph H,*

$$\lambda(G \circledz H) \geq \frac{c^2}{\sqrt{c-1}} - o_n(1).$$

For the particular case of $c = 2$ we get the bound $\lambda(G \circledz H) \geq \frac{3\sqrt{3}}{\sqrt{2}} - o_n(1) \approx 3.674$. Theorem 2.2 provides a definite answer to the previously posed question: given that the degree of $G \circledz H$ is $D = c^2$, the theorem asserts that $\omega(G \circledz H) = \Omega(D^{-\frac{1}{4}})$, demonstrating that the classical analysis provided in [RVW00] is asymptotically tight, at least for vertex-transitive graphs. Returning to the question raised in Section 1.2, we observe a notable contrast between the Zig-Zag product and the derandomized square, despite their similarity: while there is a strong evidence that $G \circleds H$ can be arbitrarily close to Ramanujan—though its classical analysis, the Rozenman-Vadhan bound, fails to disclose this—the aforementioned result *provably* indicates that the Zig-Zag product is intrinsically limited, and its capabilities are accounted for by its classical analysis. We provide a proof sketch of Theorem 2.2 in Section 3.2.

2.2 Matching the bound, and connection to free probability theory

As discussed above, we prove that the lower bound given by Theorem 2.1 is optimal, at least in a one-sided manner, by showing the existence of graphs *of every size* matching it in terms of the second largest eigenvalue.

Theorem 2.3. *For every good (d,c)-graph H and for every integer $n \geq 1$, there exists an (n,d)-graph G such that*

$$\lambda_2(G \circledz H) \leq \min_{x > c^2} \mathcal{Z}_H(x).$$

Interestingly, unlike Theorem 2.1, the upper bound of Theorem 2.3 applies to graphs H which are not necessarily vertex-transitive. Note also that we are only able to bound the second-largest eigenvalue, λ_2, rather than the spectral expansion $\lambda = \max(\lambda_2, |\lambda_n|)$. Graphs characterized by a bounded second-largest eigenvalue are called *one-sided* spectral expanders. These are also highly applicable in various contexts, mainly because they satisfy the conditions of the Alon-Chung Lemma [AC88]. The proof of Theorem 2.3 leverages finite free probability and the interlacing technique [MSS15, MSS18, MSS22]. We provide an overview for the proof of Theorem 2.3 in Section 3.3. It remains open to explicitly construct a graph G achieving the bound given by Theorem 2.3, analogous to the explicit construction of Cohen [Coh16] for bipartite Ramanujan graphs.

It is important to highlight that the alignment of our lower bound, which is based on analytic combinatorics, with our upper bound, derived from free probability theory, points to a deep connection between these two fields, as previously observed in [CCMP23] in the context of derandomized squaring. In this paper, we delve further into this connection, where we diverge from prior work by leveraging it in our proofs. Specifically, during the process of establishing our lower bound, we encounter the challenge of solving a rather complicated equation that emerges from results in analytic combinatorics. A crucial technical innovation we introduce is the application of free probability theory—particularly, insights gained from the proof of our *upper* bound—to address this otherwise daunting equation, thus establishing our *lower* bound. This is, in fact, more apparent in our analysis of the replacement product, which we discuss next.

2.3 Tight bounds for the spectral expansion of the replacement product

In analogous fashion to our lower bound and matching upper bound for the Zig-Zag product, we provide bounds for the even more fundamental replacement product $G \circledr H$ (see Definition 4.3). For brevity, we state both bounds in the theorem below, noting it is more convenient to state our result, and in particular the function $\mathcal{R}_H(x)$ which is the analog of $\mathcal{Z}_H(x)$, in terms of the Cauchy transform.

Theorem 2.4. *Let H be a good*[3] *(d,c)-graph, and let $\Lambda_H = \min_{x>c} \mathcal{R}_H(x)$, where*

$$\mathcal{R}_H(x) = x + \frac{\sqrt{1+4\mathcal{G}_H(x)^2}-1}{2\mathcal{G}_H(x)}.$$

Then, the following hold:

1. *For every (n,d)-graph G, $\lambda(G ⓡ H) \geq \Lambda_H - o_n(1)$.*
2. *For every n there exists an (n,d)-graph G such that $\lambda_2(G ⓡ H) \leq \Lambda_H$.*

The upper and lower bounds are expressed and proven separately in Theorems 6.1 and 6.3, respectively. The upper bound is straightforward given the free probability based proof of the Zig-Zag product, noticing that their matrix representations are similar, with multiplication being substituted by addition. Intriguingly, the analytic tools required for proving the lower bound in the context of the replacement product are substantially more sophisticated, despite the simplicity of the operation from the combinatorial perspective.

The replacement product, on its own, is not specifically designed to produce good expanders. Nonetheless, it is valuable to establish a universal bound for this operation, applicable to all vertex-transitive degree-c graphs, similar to Theorem 2.2. We prove that the best possible spectral expansion for the replacement product involving such c-regular graphs is given by $\lambda(G ⓡ H) \geq c + \frac{1}{c-1} - o_n(1)$, or $\omega(G ⓡ H) \geq \frac{c}{c+1} - o_n(1)$.

3 Proof overview

In this section, we provide an informal overview of the proofs for our results. In Section 3.1, we discuss our lower bound for the spectral expansion of the Zig-Zag product, as given by Theorem 2.1. Our proof relies on the symbolic method and leverages results from analytic combinatorics, both of which we introduce and explain in the respective sections (see Section 3.1 as well as Appendix A for the necessary background). With this, the proof of Theorem 2.1 is outlined in Section 3.2. We follow with an outline of the ideas for proving the analog bound for the replacement product. Additionally, we outline the proof of tightness of our lower bounds, as stated in Theorem 2.3, in Section 3.3. In that section, we provide the necessary background on finite free probability.

[3] In this context, the definition of "good" differs from the one we previously established for the Zig-Zag product; refer to Section 6.2 for details.

3.1 Analytic combinatorics and the symbolic method

Our lower bound results are based on the symbolic method which provides a framework to convert a specification of a so-called combinatorial class by means of certain combinatorial constructs into a functional equation that is satisfied by its associated generating function. Here we briefly cover only the parts that we need from this theory, where more information is given in Appendix A. The interested reader is referred to the excellent book by Flajolet and Sedgewick [FS09] to learn more about this fascinating topic.

A *combinatorial class* $\mathcal{A}$ consists of a collection of objects paired with a designated *size* function $|\cdot| : \mathcal{A} \to \mathbb{N}$. The associated *generating function* for this class is the formal power series

$$A(z) = \sum_{a \in \mathcal{A}} z^{|a|} = \sum_{k \in \mathbb{N}} A_k z^k,$$

where A_k, which is also denoted as $[z^k]A(z)$, is the number of objects in $\mathcal{A}$ of size k, which we always assume is finite. Set theoretic operators on the combinatorial classes reflect in their associated generating functions. For instance, the *sequence* of a class, denoted as $\mathsf{SEQ}(\mathcal{A})$ represents the disjoint union of the Cartesian products across all finite lengths $n \geq 0$. The generating function for $\mathsf{SEQ}(\mathcal{A})$ is given by $\frac{1}{1-A(z)}$. The elements of size 1 in a combinatorial class $\mathcal{A}$ are called *atoms*, all of which are considered distinct.

The symbolic method classifies combinatorial classes into *schemas* based on their shared structures. This approach aims to consolidate solutions to these problems and highlight their interrelations. A notable schema within this framework is termed *smooth inverse-function schema*. These are classes whose generating function $\zeta(z)$ satisfies the functional equation $u = z \cdot \phi(u)$, namely, $\zeta(z) = z \cdot \phi(\zeta(z))$, for some "well-behaved" complex function $\phi(u)$. Analytic combinatorics provides a method to estimate the coefficients of the generating function for smooth inverse-function schema. This approach is applicable under certain technical conditions on $\phi(u)$, which we hide under the rug in this informal proof overview. The key requirement though is that there is a unique real positive solution to the *characteristic equation* $\phi(u) = u \cdot \phi'(u)$ within ϕ-s analytic domain around the origin. With this, we have the following theorem which is informally stated here.

Theorem 3.1. *Let $\zeta(z)$ belong to the smooth inverse-function schema. Then, with τ the unique positive root of the corresponding characteristic equation $\phi(u) = u \cdot \phi'(u)$, one has*

$$[z^n]\zeta(z) \approx \phi'(\tau)^n. \tag{6}$$

A more general schema is the *smooth implicit-function* schema. It generalizes the above result to classes whose generating functions $\zeta(z)$ satisfy a more complicated functional equation $w = P(z, w)$, namely $\zeta(z) = P(z, \zeta(z))$ for some "well behaved" bivariate complex function $P(z, w)$. Analytic combinatorics provides a method to estimate the coefficients of the generating function also for the smooth implicit-function schema. We require these more sophisticated tools for the analysis of the replacement product though we choose not to delve into this in this section.

3.2 Lower bounding the spectral expansion of Zig-Zag products

In this section, we outline the key elements of the proof of Theorem 2.1. Let G be a d-regular graph and H a vertex-transitive (d, c)-graph. Our starting point is standard, relying on the trace method which asserts that $\lambda(G \textcircled{z} H)$ is lower bounded by roughly $c_\ell(G \textcircled{z} H)^{1/\ell}$ for every $\ell > 0$, where $c_\ell(G \textcircled{z} H)$ is the number of length-ℓ cycles that originate at some fixed vertex v of $G \textcircled{z} H$ (see Lemma 4.5). Thus, the task at hand is to lower bound $c_\ell(G \textcircled{z} H)$, where we will eventually choose ℓ to be sufficiently large. A common strategy for this is to consider a suitable infinite cover [4] of the graph of interest, $G \textcircled{z} H$ in our case, which we take to be $\mathcal{T}_d \textcircled{z} H$, where $\mathcal{T}_d$ is the d-regular infinite tree. Indeed, for every ℓ, every length-ℓ cycle in $\mathcal{T}_d \textcircled{z} H$ that originated at some fixed vertex induces a unique cycle in $G \textcircled{z} H$, initiated at some corresponding vertex, hence $c_\ell(G \textcircled{z} H) \geq c_\ell(\mathcal{T}_d \textcircled{z} H)$.

In order to apply the symbolic method more easily and specify some recursive relation on the class of cycles, we truncate some of the edges of $\mathcal{T}_d \textcircled{z} H$. We fix an arbitrary root r in $\mathcal{T}_d \textcircled{z} H$ and remove some of the edges adjacent to r, resulting in a new graph which we call X. See Figure 1 for an illustration of the truncation we use. As removing edges can only reduce the number of cycles, it suffices to lower bound $c_\ell(X)$. To this end, we specify a combinatorial class $\mathcal{S}_X$ which is the class of cycles in X, originating in the root r that only visit r upon completing a cycle. We identify an isomorphism between the restriction of the cycles in $\mathcal{S}_X$ to a single copy of H in $\mathcal{T}_d \textcircled{z} H$ to certain cycles in the graph H^2, captured by the class we denote $\mathcal{A}_H$. In Claim 5.3 we obtain an analytic formulation of $\mathcal{A}_H$-s generating function, denoted $A_H(z)$, in terms of the Cauchy transform of H^2. Particularly, we deduce that the radius of convergence of $A_H(z)$ is $\frac{1}{c^2}$. We then establish the following recursive symbolic relation on $\mathcal{S}_X$:

$$\mathcal{S}_X = \mathcal{Z} \times (\mathcal{A}_H \circ \mathcal{S}_X) \times \mathcal{Z},$$

where $\mathcal{Z}$ is an atom (see Lemma 5.4). From the symbolic relation, we immediately derive a functional equation that is satisfied by the corresponding generating function,

$$S_X(z) = z^2 \cdot A_H(S_X(z)).$$

Recall that $[z^\ell]\, S_X(z) \geq c_\ell(X)$. Thus, with the functional equation in hand, our objective is to deduce estimates of its coefficients using Theorem 3.1. For invoking the theorem, we define a function $E(z)$ such that $E(z^2) = S_X(z)$ and such that $[z^{2n}]S_X(z) = [z^n]E(z)$, which

[4] We will not need the formal definition of a cover in this paper, and we mention it here by name for the reader who is familiar with this notion.

satisfies the functional equation $E(z) = z \cdot A_H(E(z))$. We then prove in Claim 5.11 that $E(z)$ belongs to the smooth inverse-function schema. All the technical conditions are easily established therein. In Lemma 5.8, we show that the key component, the characteristic equation $A_H(u) = u \cdot A'_H(u)$ is equivalent, up to a change of variable $u = \frac{1}{x}$, to the equation $\mathcal{Z}'_H(x) = 0$, where $\mathcal{Z}_H(x)$ is given by Equation (3).

We prove in Section 5.1.2 that for every good H, the function $\mathcal{Z}_H(x)$ must attain a unique minimum x_0 in the range (c^2, ∞). This minimum then induces the solution τ to the characteristic equation $A_H(u) = u \cdot A'_H(u)$ within its radius of convergence. We further prove in Section 5.1.2 that $\sqrt{A'_H(\tau)} = \mathcal{Z}_H(x_0)$, which by Theorem 3.1, determines the asymptotic growth of the coefficients of $E(z)$ and hence also of $S_X(z)$. Putting everything together completes the proof.

Zig-Zag is inherently quadratic far from Ramanujan. Equipped with Theorem 2.1, we turn to discuss our lower bound that holds for all good c-regular vertex-transitive graphs H (the proof appears in Section 5.2). As mentioned, the intuition behind the proof is that $\mathcal{T}_c$, the infinite c-regular tree, is a universal cover for all c-regular graphs, and therefore has at most as many cycles of any size. Hence, the Cauchy transform $\mathcal{G}_{\mathcal{T}_c^2}$, appropriately defined, could be plugged into the definition of $\mathcal{Z}_H$ to achieve a universal lower bound.

More formally, we observe that the function $\Psi(x, y) = \sqrt{x^2 - \frac{x}{y}}$ is monotone increasing in y. Taking a closer look at $\mathcal{Z}_H(x)$, the dependence in H is captured in the variable y of $\Psi(x, y)$. Since for *any* (d, c)-graph H and for all $x > c^2$ we have that $\mathcal{G}_{H^2}(x) \geq \mathcal{G}_{\mathcal{T}_c^2}(x)$, we get that $\mathcal{Z}_H(x) \geq \mathcal{Z}_{\mathcal{T}_c}(x)$ for all $x > c^2$. The proof follows by observing that for $c \geq 3$,

$$\min_{x>c^2} \mathcal{Z}_H(x) \geq \inf_{x>c^2} \mathcal{Z}_{\mathcal{T}_c}(x) = \mathcal{Z}_{\mathcal{T}_c}(c^2) = \frac{c^2}{\sqrt{c-1}}.$$

3.3 Existence of graphs matching the bound

To match the lower bound described in Section 3.2, given a fixed (d, c)-graph H and any integer $n \geq 1$, we wish to find an (n, d)-graph G such that $\lambda_2(G \textcircled{z} H) \leq \min_{x>c^2} \mathcal{Z}_H(x)$, as stated in Theorem 2.3. First, observe that the adjacency matrix of $G \textcircled{z} H$ can be expressed as

$$\mathbf{Z} = \boldsymbol{\mathcal{H}} \dot{\mathbf{G}} \boldsymbol{\mathcal{H}}, \tag{7}$$

where $\boldsymbol{\mathcal{H}} = \mathbf{I}_n \otimes \mathbf{A}_H$, $\mathbf{A}_H$ being the adjacency matrix of H, and $\dot{\mathbf{G}}$ is the matrix which encodes the rotation map of G, dubbed the *rotation matrix* of G. Therefore, we have reformulated our problem as finding a graph G where the second largest eigenvalue of the matrix $\mathbf{Z}$ is small. This rephrasing opens the door to applying *finite free probability*. To begin, in Section 3.3.1, we provide a brief overview of this elegant theory, highlighting its relevance to the proof of Theorem 2.3. It is important to note that the discussion in this high-level summary is in fact more detailed than the actual proof presented in Section 5.3, as this section delves into the methodology underpinning the proof.

3.3.1 A brief introduction to finite free probability

Free probability is a branch of mathematics, initiated by Voiculescu, that extends classical probability theory into the non-commutative setting. In classical probability, random variables are analyzed using their joint distribution, which encodes the correlations or lack of between them. In contrast, free probability introduces the abstract notion of "freeness" to represent the absence of correlations, appropriately defined, among non-commutative random variables. We invite the reader to learn more about this theory in the introductory book by Nica and Speicher [NS06].

Free probability theory provides, in particular, tools to analyze the spectrum of the sum and product of two operators, given that they are free, using knowledge of their individual spectra. For finite matrices, though, the theory provides mainly asymptotic results, where the dimension of the matrices tends to infinity. As a result, operators associated with finite graphs cannot be studied directly by free probability theory.

In response to this limitation, Marcus, Spielman, and Srivastava, in their groundbreaking series of works [MSS15, MSS18, MSS22], introduced the theory of *finite free probability* along with the associated technique of interlacing. This enabled them to extend some results of free probability to the finite-dimensional setting, especially regarding the spectra of matrix sums and products. [5] The *finite free convolutions* defined in [MSS22], presented below, can be defined in a standalone fashion, not relying on the abstract concept of freeness. Rather, they demonstrate that conjugating a finite operator, $\mathbf{A}$, with a Haar-orthogonal matrix, effectively "frees" $\mathbf{A}$ from other operators. We consider a specific application of this principle in the context of operator multiplication, which will turn useful in analyzing the spectrum of $\mathbf{Z}$ from Equation (7), when the graph G is picked at random.

3.3.2 Existence of graphs matching the Zig-Zag bound

Before turning to the formal definitions used in the proof, we outline the basic idea of the randomized method which underlies the proof of Theorem 2.3. This involves three steps:

1. Define a distribution by which the graph G is drawn.

2. Analyze the *expected characteristic polynomial* of the Zig-Zag product of the random G and a fixed graph H, using free probability theory.

3. Use the technique of *interlacing* to conclude that there exists a specific graph G in the support of the distribution satisfying a desired property of the above expected polynomial.

[5] Already here we wish to stress, that the replacement and Zig-Zag operations will correspond, respectively, to the sum and product of a random matrix with some fixed matrix.

Note that the above steps, although formalized (and analyzed) separately, are not independent: one needs to be attentive in the choice of a distribution in Item 1 that it enables Items 2 and 3. In our analysis, diverging from prior works [MSS18, CCMP23], this distribution is taken according to the *configuration model*, which enables us to express the rotation map of G as $\dot{\mathbf{G}} = \mathbf{PMP}^\intercal$, where $\mathbf{M}$ is the adjacency matrix of an arbitrary fixed perfect matching of dimension nd, and $\mathbf{P}$ is a random permutation matrix of the same dimension (see Section 4.1 for the relevant standard definitions). This enables us to express Equation (7) as

$$\mathbf{Z}_\mathbf{P} = \boldsymbol{\mathcal{H}}\mathbf{PMP}^\intercal\boldsymbol{\mathcal{H}}, \tag{8}$$

and the expected characteristic polynomial we wish to analyze is then

$$\mathop{\mathbf{E}}_{\mathbf{P}} \chi_x(\mathbf{Z}_\mathbf{P}) = \mathop{\mathbf{E}}_{\mathbf{P}} \chi_x\left(\boldsymbol{\mathcal{H}}\mathbf{PMP}^\intercal\boldsymbol{\mathcal{H}}\right) = \mathop{\mathbf{E}}_{\mathbf{P}} \chi_x\left(\boldsymbol{\mathcal{H}}^2\mathbf{PMP}^\intercal\right), \tag{9}$$

where the last equality is due to the invariance of the characteristic polynomial over cyclic rotations of the matrices. This sets the stage for the definitions relevant for analyzing this polynomial in Item 2.

Definition 3.2 (Free multiplicative convolution)**.** *Let* $\mathbf{A}, \mathbf{B}$ *be* $m \times m$ *real symmetric matrices, with characteristic polynomials* $a(x)$ *and* $b(x)$*, respectively. The* free multiplicative convolution *of* $a(x)$ *and* $b(x)$ *is defined as*

$$a(x) \boxtimes b(x) = \mathop{\mathbf{E}}_{\mathbf{Q}} \chi_x\left(\mathbf{AQBQ}^\intercal\right), \tag{10}$$

where the expectation is taken in the coefficient space over random orthogonal matrices $\mathbf{Q}$ *sampled according to the Haar measure on the group of* n*-dimensional orthogonal matrices.*

There is a clear resemblance between Equation (10) and the right hand side of Equation (9), where $\boldsymbol{\mathcal{H}}^2$ plays the role of $\mathbf{A}$, and $\mathbf{M}$ plays the role of $\mathbf{B}$. The main difference is our wish to analyze an expectation over *permutation* matrices, while Definition 3.2 deals with Haar orthogonal matrices. Following MSS, we are able to overcome this obstacle using a *quadrature* argument, formalized in Lemma 4.8.

To recap, for bounding the roots of the expected characteristic polynomial $\mathbf{E}_\mathbf{P}\,\chi_x(\mathbf{Z}_\mathbf{P})$, we are left with the task of bounding those of $(p_{\boldsymbol{\mathcal{H}}^2} \boxtimes p_\mathbf{M})(x)$, where $p_{\boldsymbol{\mathcal{H}}^2}(x)$ and $p_\mathbf{M}(x)$ are essentially the characteristic polynomials of the respective matrices. The analytic framework that will allow us to study the multiplicative convolution includes the Cauchy transform and the moment transform, and their respective inverse functions. Recall from Equation (5) that the Cauchy transform of a real-rooted degree d polynomial $p(t) \in \mathbb{R}[t]$, whose roots are $\lambda_1 \geq \lambda_2 \geq \cdots \geq \lambda_d$, is defined as $\mathcal{G}_p(x) = \frac{1}{d}\sum_{i=1}^{d} \frac{1}{x-\lambda_i}$. For an undirected graph H, we write $\mathcal{G}_H(x)$ for the Cauchy transform $\mathcal{G}_p(x)$ where $p(t) = \chi_t(H)$ is the characteristic polynomial of H. The $\mathcal{M}$-transform (or moment transform), is defined by

$\mathcal{M}_p(x) = x\mathcal{G}_p(x) - 1$. We define the inverse of $\mathcal{M}_p(x)$, denoted $\mathcal{N}_p(y)$, to be the largest x so that $\mathcal{M}_p(x) = y$, assuring the reader that this function is well-defined. It is worth noting that by definition, for every $y \in (0, \infty)$, the value $\mathcal{N}_p(y)$ provides an upper bound on λ_1. The key feature of the $\mathcal{N}$-transform in our context is that it behaves well under free multiplicative convolutions. We state here the theorem encapsulating this idea, which appears in [MSS22].

Theorem 3.3. *Let p and q be polynomials of the same degree, of which at least one has only non-negative roots*[6]*. Then, for every $y > 0$,*

$$\mathcal{N}_{p \boxtimes q}(y) \leq \frac{y}{y+1} \cdot \mathcal{N}_p(y)\mathcal{N}_q(y). \tag{11}$$

At this point we see that $\mathcal{N}_{\mathbf{M}}(y)$ can be calculated explicitly, and $\mathcal{N}_{\boldsymbol{\mathcal{H}}^2}(y)$ can be expressed in terms of the Cauchy transform of H^2. Therefore, applying Theorem 3.3 to Equation (9) along with the quadrature technique (and a few straightforward calculations) yields

$$\alpha_2 \left(\mathop{\mathbf{E}}_{\mathbf{P}} \chi_x \left(\mathbf{Z}_{\mathbf{P}} \right) \right) \leq \min_{x > c^2} \sqrt{x^2 - \frac{x}{\mathcal{G}_{H^2}(x)}} = \min_{x > c^2} \mathcal{Z}_H(x),$$

where $\alpha_k(p)$ is the k-th largest root of the polynomial $p(x)$. As we have bounded the roots of the expected characteristic polynomial, we turn to Item 3, which is the deduction of the existence of a graph having similar properties to it.

Interlacing. So far, we have discussed how to obtain a bound on the largest root of the expected characteristic polynomial, where the expectation is in coefficient space and taken over the group of permutation matrices. It is generally incorrect to assert that a bound on the largest root of the expectation of polynomials can be utilized to infer a bound on the largest root of one of the polynomials involved in the expectation. A key observation by MSS concerning this issue is that such a result holds if the polynomials participating in the expectation form an *interlacing family*, a condition satisfied in our case. This idea is formalized in Lemma 4.9, where it is shown that in fact, for any choice of k, this structure suffices to deduce a bound on the k-th largest root of at least one polynomial in the family, given that we are able to bound the k-th largest root of the expected characteristic polynomial.

3.3.3 Existence of graphs matching the replacement product bound

Our upper bound proof for the replacement product follows the same argument discussed in Section 3.3.2, roughly by replacing multiplication with addition. The adjacency matrix of the replacement product $G \circledR H$ can be expressed as $\mathbf{R} = \boldsymbol{\mathcal{H}} + \dot{\mathbf{G}}$, similarly to Equation (7).

[6]Originally, [MSS22] proved the theorem when both polynomials have non-negative roots. However, we note that their proof generalizes to this case, as we further discuss in Section 4.2.1.

Representing G in the configuration model enables us to express this equation as $\mathbf{R_P} = \boldsymbol{\mathcal{H}} + \mathbf{PMP}^{\mathsf{T}}$, similarly to Equation (8). As the $\mathcal{N}$ transform and free multiplicative convolution served us in proving an upper bound on the roots of $\mathbf{E_P}\,\chi_x(\mathbf{Z_P})$, the exact same strategy using the $\mathcal{K}$ transform (which is the inverse of the Cauchy transform $\mathcal{G}$) and the free *additive* convolution lead us to a bound on the spectrum of the expected characteristic polynomial $\mathbf{E_P}\,\chi_x(\mathbf{R_P})$.

4 Preliminaries

In this section we set notation and shortly survey the necessary background for our proofs: graphs and graph products in Section 4.1, the Cauchy-transform and finite free convolutions in Section 4.2, and finally analytic combinatorics in Section 4.3.

4.1 Graphs

4.1.1 Rotation maps and the configuration model

Let G be an (n, d)-graph, and assume that every vertex v has a labeling of the d edges adjacent to it with the labels $\{1, \dots, d\}$ such that every label appears exactly once. Notice that each edge is labeled twice (by both its end vertices), and the two labels might differ. Let $v[i]$ denote the i-th neighbor of v according to this labeling.

Definition 4.1 (Edge rotation map)**.** *The* edge rotation map *(or simply* rotation map*) of a labeled (n, d)-graph $G = (V, E)$, denoted $\mathrm{Rot}_G : V \times [d] \to V \times [d]$, is defined by*

$$\mathrm{Rot}_G(v, i) = (u, j) \iff v[i] = u \;\wedge\; u[j] = v.$$

Otherwise put, $\mathrm{Rot}_G(v, i) = (u, j)$ if the i-th neighbor of v is u, and the j-th neighbor of u is v. Notice that Rot_G is an involution. We say that G is *consistently labeled* if for every edge $\{u, v\} \in E$, $\mathrm{Rot}_G(v, i) = (u, i)$ for some $i \in [d]$; that is, every edge is seen with the same label from both its end points. We also define accordingly the *rotation matrix* $\dot{\mathbf{G}}$, which is an $nd \times nd$ boolean matrix, where $\dot{\mathbf{G}}_{(v,i),(u,j)} = 1$ if and only if $\mathrm{Rot}_G(v, i) = (u, j)$.

We proceed by defining the *configuration model* for sampling a random (n, d)-graph, for $n, d \in \mathbb{N}$ of which at least one is even. To this end, it might be helpful to envision an initial phase in which each vertex is connected to d "half-edges", and the graph is specified by connecting the nd half edges.

Definition 4.2 (Configuration model)**.** *A random (n, d)-graph is said to be sampled by the* configuration model *if the distribution can be described by sampling a uniformly random perfect matching of the nd* half-edges*.*

The model can also be described, more formally, in matrix form: let $\mathbf{M}$ be the adjacency matrix of an arbitrary perfect matching of dimension nd. Then, picking G according to the configuration model is equivalent to picking a random $nd \times nd$ permutation matrix $\mathbf{P}$ and setting $\dot{\mathbf{G}} = \mathbf{P}\mathbf{M}\mathbf{P}^{\intercal}$. We remark that unlike some other random graph models, self-loops are allowed: it is possible that $\mathrm{Rot}_G(v,i) = (v,j)$ for some vertex v and $i,j \in [d]$, in which case v is its own i-th and j-th neighbor. We proceed by defining the two graph products which will be of interest in our main results.

4.1.2 The replacement and Zig-Zag products

Definition 4.3 (Replacement product)**.** *Let $G = (V_G, E_G)$ be an (n,d)-graph and $H = (V_H, E_H)$ a (d,c)-graph. The graph $G \circledR H = (V_G \times V_H, E)$ is an $(nd, c+1)$-graph defined as follows: for every $v \in V_G$, $\{(v,i),(v,j)\} \in E$ for every $\{i,j\} \in E_H$. In addition, $\{(v,i),(u,j)\} \in E$ whenever $\mathrm{Rot}_G(v,i) = (u,j)$.*

One can think of the replacement product in the following way: Every vertex v of G is replaced by a copy of H which we dub the *cloud* of v, where the original edges from G are connecting the clouds according to Rot_G. Observe that the adjacency matrix of $G \circledR H$ can be written as

$$\mathbf{R} = \boldsymbol{\mathcal{H}} + \dot{\mathbf{G}},$$

where $\boldsymbol{\mathcal{H}} = \mathbf{I}_n \otimes \mathbf{A}_H$, and $\mathbf{A}_H$ is the adjacency matrix of H.

Definition 4.4 (Zig-Zag product [RVW00])**.** *Let $G = (V_G, E_G)$ be an (n,d)-graph and $H = (V_H, E_H)$ a (d,c)-graph. The graph $G \circledz H = (V_G \times V_H, E)$ is the (nd, c^2)-graph defined as follows: the edge $\{(v,i),(u,j)\}$ is added to E for every ℓ_1, ℓ_2 such that $\mathrm{Rot}_G(v,\ell_1) = (u,\ell_2)$, and $\{i,\ell_1\}, \{j,\ell_2\}$ are edges in E_H.*

In a similar manner to the replacement product, we think of a cloud replacing every $v \in V_G$, where in this case the edges can be interepreted by the following walk: from (v,i) one first takes a step in H within the same cloud. Then, uses Rot_G to transition to another cloud, and finally takes an additional H step in the latter[7]. Accodring to Definition 4.4, the adjacency matrix of $G \circledz H$ can be written as

$$\mathbf{Z} = \boldsymbol{\mathcal{H}}\dot{\mathbf{G}}\boldsymbol{\mathcal{H}}.$$

[7]For defining the rotation map of the Zig-Zag product, one can write $\mathrm{Rot}_{G \circledz H}((v,i),(a_1,a_2)) = ((u,j),(b_1,b_2))$ if, using the notations above, $\mathrm{Rot}_G(v,\ell_1) = (u,j)$, $\mathrm{Rot}_H(i,a_1) = (\ell_1, b_2)$ and $\mathrm{Rot}_H(j,a_2) = (\ell_2, b_1)$. This definition is necessary for using the Zig-Zag operation in recursive constructions [RVW00], but will not be needed for our analysis.

4.1.3 The trace method

In this paper we use a very common technique for lower bounding spectral properties of the graph, called the *trace method*. The merit of this method is reducing the task of bounding the spectral expansion of a graph to the task of lower bounding the number of cycles originating at any vertex. Denote by $C_k(G, v)$ the number of cycles of length k, originating at a vertex v in the graph G. The following lemma appears in various contexts, and is considered a folklore. For a proof see, e.g., Lemma 4.2 in [CCMP23].

Lemma 4.5 (Trace method for regular graphs)**.** *Let F be a family of graphs such that for every $G = (V, E) \in F$ and $v \in V$ we have that*

$$(C_{2k}(G, v))^{\frac{1}{2k}} \geq (1 - o_k(1))\, \rho,$$

for some constant $\rho > 0$ that is independent of G and for infinitely many k-s. Then, for every $G \in F$,

$$\lambda(G) \geq (1 - o_n(1))\, \rho, \tag{12}$$

where n is the number of vertices of the graph G.

It is common to prove a lower bound on $C_{2k}(G, v)$ by looking at a cover for all graphs in F. For example, the most famous and important lower bound on the spectral expansion of regular graphs is the aforementioned Alon-Boppana bound, which can be proved using the trace method, analyzing the d-regular infinite tree $\mathcal{T}_d$ which is the universal cover of all d-regular graphs.

4.2 Finite free probability

We use the following standard notation: for a symmetric matrix $\mathbf{A}$, $\chi_x(\mathbf{A})$ is the characteristic polynomial of $\mathbf{A}$ with variable x. We denote by $\lambda_k(\mathbf{A})$ the k-th largest eigenvalue of $\mathbf{A}$. For a real-rooted polynomial $p(x)$, we denote by $\alpha_k(p)$ the k-th largest root of $p(x)$. For a distribution P over polynomials, we denote by $\mathbf{E}_{p\sim P}[p(x)]$ the *expected* polynomial over this distribution, where the expectation is taken in coefficient space, namely, for every k, the coefficient of x^k in $\mathbf{E}_{p\sim P}[p(x)]$ is the expectation over coefficients corresponding to x^k in $p(x)$ drawn according to P.

Definition 4.6 (Haar distribution on the orthogonal group)**.** *Denote the group of $m \times m$ orthogonal matrices by $\mathcal{O}(m)$. The* Haar distribution *is the unique distribution over $\mathcal{O}(m)$ which is invariant under multiplication (from the right or from the left) with any orthogonal matrix. We call a matrix drawn from this distribution a* Haar random matrix*.*

An important characteristic of the Haar distribution, upon which Definition 4.7 below relies, is the following. Let $\mathbf{A}, \mathbf{B}$ be two arbitrary $m \times m$ symmetric matrices, and $\mathbf{Q}$

a Haar random matrix of the same dimensions. Informally, the random rotation of $\mathbf{B}$ according to $\mathbf{Q}$ removes any dependence between the respective eigenvectors of $\mathbf{A}$ and $\mathbf{B}$. More formally, if $\chi_x(\mathbf{A}) = a(x)$ and $\chi_x(\mathbf{B}) = b(x)$, then both expected characteristic polynomials $\mathbf{E}_{\mathbf{Q}}\, \chi_x(\mathbf{A} + \mathbf{Q}\mathbf{B}\mathbf{Q}^\intercal)$ and $\mathbf{E}_{\mathbf{Q}}\, \chi_x(\mathbf{A}\mathbf{Q}\mathbf{B}\mathbf{Q}^\intercal)$ depend only on $a(x)$ and $b(x)$, and not on the eigenvectors of either $\mathbf{A}$ or $\mathbf{B}$.

Definition 4.7 (Additive and multiplicative convolutions)**.** *Let* $\mathbf{A}, \mathbf{B}$ *be real symmetric matrices of equal dimension, with characteristic polynomials* $a(x) = \chi_x(\mathbf{A})$ *and* $b(x) = \chi_x(\mathbf{B})$*. The* additive convolution $a \boxplus b$ *and the* multiplicative convolution $a \boxtimes b$ *are the polynomials defined by*

$$(a \boxplus b)(x) = \mathop{\mathbf{E}}_{\mathbf{Q}} \chi_x(\mathbf{A} + \mathbf{Q}\mathbf{B}\mathbf{Q}^\intercal),$$

and

$$(a \boxtimes b)(x) = \mathop{\mathbf{E}}_{\mathbf{Q}} \chi_x(\mathbf{A}\mathbf{Q}\mathbf{B}\mathbf{Q}^\intercal),$$

where $\mathbf{Q}$ *is a Haar random orthogonal matrix.*

Although Definition 4.7 involves the matrices $\mathbf{A}$ and $\mathbf{B}$, it depends in fact only on $a(x)$ and $b(x)$, due to the properties of the Haar measure.[8] It is important to note that, as proven in [MSS22], both $(a \boxplus b)(x)$ and $(a \boxtimes b)(x)$ are real-rooted. Interestingly, there are explicit formulas for both $(a \boxplus b)(x)$ and $(a \boxtimes b)(x)$ as functions of the coefficients of $a(x)$ and $b(x)$. We refer the reader to [MSS22] for more details.

Working with graph matrices, an issue with the above defintion is that the Haar measure does not have a meaningful combinatorial interpretation. Therefore, we need a way to relate *permutations* matrices - which do have such interpretation - to Haar random matrices. To this end we state the following lemma (which appeared first in [MSS18] for the additive case, and later on in [CM23] for the multiplicative case, whose proof uses similar ideas) which is described as a *Quadrature* result, translating an infinite (continuous) measure space to a finite one.

Lemma 4.8 (*Quadrature*; Corollary 4.9 from [MSS18] and Lemma 2.3 from [CM23])**.** *Let* $\mathbf{A}, \mathbf{B}$ *be real* $m \times m$ *symmetric matrices such that* $\mathbf{A}\mathbf{1} = a\mathbf{1}$ *and* $\mathbf{B}\mathbf{1} = b\mathbf{1}$*. Denote by* $p_{\mathbf{A}}, p_{\mathbf{B}}$ *the polynomials satisfying* $\chi_x(\mathbf{A}) = (x-a)p_{\mathbf{A}}(x)$*,* $\chi_x(\mathbf{B}) = (x-b)p_{\mathbf{B}}(x)$*. Let* $\mathbf{P}$ *be a uniformly random* $m \times m$ *permutation matrix. Then,*

$$\mathop{\mathbf{E}}_{\mathbf{P}} \chi_x\left(\mathbf{A} + \mathbf{P}\mathbf{B}\mathbf{P}^\intercal\right) = \left(x - (a+b)\right)\left(p_{\mathbf{A}} \boxplus p_{\mathbf{B}}\right)(x), \tag{13}$$

$$\mathop{\mathbf{E}}_{\mathbf{P}} \chi_x\left(\mathbf{A}\mathbf{P}\mathbf{B}\mathbf{P}^\intercal\right) = (x - ab)\left(p_{\mathbf{A}} \boxtimes p_{\mathbf{B}}\right)(x). \tag{14}$$

Another tool we will need is *Interlacing*, which in our context will enable us deduce

[8] It is easily seen that the convolution is well defined for any real-rooted polynomials $a(x)$ and $b(x)$ by choosing $\mathbf{A}, \mathbf{B}$ to be diagonal matrices with their respective roots on the diagonal.

bounds on roots of specific polynomials using similar bounds on the roots of an expectation polynomial.

Lemma 4.9 (*Interlacing*). *Suppose* $\mathbf{A}, \mathbf{B}$ *are symmetric* $m \times m$ *matrices and* $\mathbf{P}$ *is a uniform random* $m \times m$ *permutation matrix. Then, for every* $k \leq m$ *there exist permutation matrices* $\mathbf{R}, \mathbf{S}$ *such that*

$$\lambda_k \left(\mathbf{A} + \mathbf{R}\mathbf{B}\mathbf{R}^\intercal\right) \leq \alpha_k \left(\mathop{\mathbf{E}}_{\mathbf{P}} \chi_x \left(\mathbf{A} + \mathbf{P}\mathbf{B}\mathbf{P}^\intercal\right)\right), \tag{15}$$

$$\lambda_k \left(\mathbf{A}\mathbf{S}\mathbf{B}\mathbf{S}^\intercal\right) \leq \alpha_k \left(\mathop{\mathbf{E}}_{\mathbf{P}} \chi_x \left(\mathbf{A}\mathbf{P}\mathbf{B}\mathbf{P}^\intercal\right)\right), \tag{16}$$

where we recall that $\alpha_k(p)$ *is the* k*-th largest root of* $p(x)$*.*

Lemma 4.9 is a simpler form of a much more general statement, appearing in Theorem 3.4 of [MSS18] (for Equation (15)) and in Theorem 6.5 and Lemma 6.3 of [CM23] (for Equation (16)).

4.2.1 Transforms

Let $p(x)$ be a degree m real-rooted polynomial with roots $\alpha_1 \geq \alpha_2 \geq \cdots \geq \alpha_m$. The *Cauchy transform* of $p(x)$ is defined as the function

$$\mathcal{G}_p(x) = \frac{1}{m} \sum_{i=1}^{m} \frac{1}{x - \alpha_i} = \frac{1}{m} \cdot \frac{p'(x)}{p(x)}.$$

When $p(x)$ is the characteristic polynomial of an $m \times m$ matrix $\mathbf{A}$, it holds that

$$\mathcal{G}_p(x) = \frac{1}{m} \mathsf{Tr}\left((x\mathbf{I} - \mathbf{A})^{-1}\right).$$

In many settings it is instructive to study the Cauchy transform as a function whose domain is $\mathbb{C}^+$. However, we will consider the Cauchy transform as a function on $\mathbb{R}$, where we evaluate the Cauchy transform to the right of its rightmost pole, at $x > \alpha_1$.

Note that when the Cauchy transform of a polynomial p is restricted to (α_1, ∞), its range is $(0, \infty)$. Additionally, $\mathcal{G}_p(x)$ is monotonically decreasing within this domain. With this in mind, one can define $\mathcal{K}_p : (0, \infty) \to (\alpha_1, \infty)$ as the inverse of $\mathcal{G}_p(x)$ when restricted to the latter domain. In other words, $\mathcal{K}_p$ is the max-inverse of $\mathcal{G}_p$. Particularly, for every $y \in (0, \infty)$, $\mathcal{K}_p(y)$ provides an upper bound on α_1, the largest root of p.

Accompanied to the Cauchy transform of p is the $\mathcal{M}$*-transform* (or moment transform), which is defined by

$$\mathcal{M}_p(x) = x\mathcal{G}_p(x) - 1. \tag{17}$$

In a similar manner to the definition of $\mathcal{K}_p$, we define the inverse of $\mathcal{M}_p$, denoted $\mathcal{N}_p(y)$, to be the largest x so that $\mathcal{M}_p(x) = y$. In a similar manner to $\mathcal{K}_p$, for every $y \in (0, \infty)$,

$\mathcal{N}_p(y)$ provides an upper bound on α_1.

For a symmetric matrix $\mathbf{A}$, we use the notation $\mathcal{G}_{\mathbf{A}}$ for $\mathcal{G}_{\chi_x(\mathbf{A})}$, and $\mathcal{G}_{p_{\mathbf{A}}}$ for the Cauchy transform of $p_{\mathbf{A}}$ as used in Lemma 4.8, which will be relevant for all matrices throughout this paper (that is, there will always exist a such that $\mathbf{A}\mathbf{1} = a\mathbf{1}$). We will similarly denote $\mathcal{M}_{\mathbf{A}}, \mathcal{K}_{\mathbf{A}}, \mathcal{N}_{\mathbf{A}}, \mathcal{M}_{p_{\mathbf{A}}}, \mathcal{K}_{p_{\mathbf{A}}}, \mathcal{N}_{p_{\mathbf{A}}}$. The key feature of the $\mathcal{K}$ and the $\mathcal{N}$ transforms is that they behave very well under additive and multiplicative convolutions, respectively, as shown in the following two lemmata.

Lemma 4.10 (Theorem 1.11 in [MSS22])**.** *For real-rooted polynomials $p(x)$ and $q(x)$ of the same degree, and for any $y > 0$,*

$$\mathcal{K}_{p \boxplus q}(y) \leq \mathcal{K}_p(y) + \mathcal{K}_q(y) - \frac{1}{y}.$$

Lemma 4.11 (Generalization of Theorem 1.12 in [MSS22])**.** *Let $p(x)$ and $q(x)$ be real-rooted polynomials such that $q(x)$ has non-negative roots and $p(x)$ has at least one positive root. Then, for every $y > 0$,*

$$\mathcal{N}_{p \boxtimes q}(y) \leq \frac{y}{y+1} \cdot \mathcal{N}_p(y)\mathcal{N}_q(y).$$

Lemma 4.11 appears in [MSS22], with the additional requirement that *both* polynomials have only non-negative roots. We state this generalization here, and provide the proof in Appendix C for completeness, though the proof follows the same arguments.

It is a common use case to apply Lemmas 4.8, 4.10 and 4.11 to graph matrices, and so one needs to calculate the $\mathcal{K}$ and $\mathcal{N}$ transforms for the polynomials $p_{\mathbf{A}}(x)$ and $p_{\mathbf{B}}(x)$. Analytically though, it is more convenient to use the same transforms of $\chi_x(\mathbf{A})$ and $\chi_x(\mathbf{B})$, which only differ from $p_{\mathbf{A}}(x)$ and $p_{\mathbf{B}}(x)$ by one root, respectively. Luckily, for graph matrices, the numbers a and b used in Lemma 4.8 are the largest eigenvalues of the respective matrices, typically representing the degree of the graphs. This enables us to make use of the following claim, which makes the above issue a mere technicality.

Claim 4.12. *Let $\mathbf{A}$ be the adjacency matrix of an (n, d)-graph G. Then, for every $x > d$, $\mathcal{G}_{p_{\mathbf{A}}}(x) \leq \mathcal{G}_{\mathbf{A}}(x)$ and $\mathcal{M}_{p_{\mathbf{A}}}(x) \leq \mathcal{M}_{\mathbf{A}}(x)$, and for every $y > 0$, $\mathcal{K}_{p_{\mathbf{A}}}(y) \leq \mathcal{K}_{\mathbf{A}}(y)$ and $\mathcal{N}_{p_{\mathbf{A}}}(y) \leq \mathcal{N}_{\mathbf{A}}(y)$.*

Proof. For the first part, notice that both $\mathcal{G}_{\mathbf{A}}(x)$ and $\mathcal{G}_{p_{\mathbf{A}}}(x)$ are averages of terms of the form $\frac{1}{x-\alpha_i}$, where the α_i-s are the roots of the respective polynomials, and $\mathcal{G}_{\mathbf{A}}(x)$ is averaging over one additional term compared to $\mathcal{G}_p(x)$, that term being $\frac{1}{x-d}$, which is the largest when evaluated at $x > d$. Therefore $\mathcal{G}_{p_{\mathbf{A}}}(x) \leq \mathcal{G}_{\mathbf{A}}(x)$, and by definition $\mathcal{M}_{p_{\mathbf{A}}}(x) \leq \mathcal{M}_{\mathbf{A}}(x)$ as well, with equality iff $\mathbf{A} = d\mathbf{I}$. For the latter two claims, we notice that these are the inverse functions of the above two which are monotone decreasing. $\square$

4.3 Analytic combinatorics and the symbolic method

We follow the notation of [FS09] and use the symbol $\bowtie$ for denoting the exponential order of sequences, which aligns with our analysis of spectra of graphs using the trace method. Formally, we say that a sequence of integers $c = (c_n)_{n\in\mathbb{N}}$ is *of exponential of order* K, and write $c \bowtie K$ if $\limsup_{n\to\infty} |c_n|^{\frac{1}{n}} = K$. Additionally, for a function $f(x)$ with expansion $f(x) = \sum_{n=0}^{\infty} f_n x^n$ around 0, we write $f(x) \bowtie K$ in the case that $f \bowtie K$ for $f = (f_n)_{n\in\mathbb{N}}$.

We also recall the definition of a *modulus of singularity*; We say that R is the modulus of singularity of a complex function $f(z)$ nearest to the origin if

$$R = \sup \{ r \geq 0 \mid f \text{ is analytic in } |z| < r \}.$$

R is also referred to as the *radius of convergence* of f. As noted in [FS09], Chapter IV, for $f(z)$ with non-negative coefficients (a class which includes all combinatorial generating functions), we have that

$$R = \sup \{ r \geq 0 \mid f(x) \text{ is analytic at all points } 0 \leq x < r \}. \tag{18}$$

Throughout the paper, we use the following theorem to find the exponential order of combinatorial generating functions.

Theorem 4.13 (Exponential growth formula; Theorem IV.7 in [FS09])**.** *If $f(z)$ is analytic at 0 and R is the modulus of singularity nearest to the origin, then*

$$f(z) \bowtie \frac{1}{R}.$$

Definition 4.14. *A function $y(z)$ analytic at 0, is said to* belong to the smooth inverse-function schema *if there exists a complex function $\phi(u)$, analytic at 0, such that in a neighborhood of 0, one has $y(z) = z \cdot \phi(y(z))$, and $\phi(u)$ satisfies the following conditions:*

$$\phi(u) \neq \alpha + \beta u, \quad \phi(0) \neq 0, \quad \textit{and} \quad \forall n \ [u^n]\phi(u) \geq 0, \tag{19}$$

and within its radius of convergence around 0, there exists a (necessarily unique) positive solution τ to the characteristic equation $\phi(u) = u \cdot \phi'(u)$*.*

Theorem 4.15 (Theorem VI.6 in [FS09], restated[10])**.** *Let $y(z)$ belong to the smooth inverse-function schema, $y(z) = z \cdot \phi(y(z))$ and let τ be the solution to the characteristic equation. Then $y(z) \bowtie \phi'(\tau)$.*

A more general theorem is applicable in the case that a more complicated recursive relation holds.

Definition 4.16. *Let $y(z)$ be a function analytic at 0 with expansion $y(z) = \sum_{n\geq 0} y_n z^n$, where $y_0 = 0$ and such that all coefficients are non negative. The function $y(z)$ is said to belong to the* smooth implicit-function schema *if there exists a bivariate function $P(z,w)$ such that*

$$y(z) = P(z, y(z)),$$

where $P(z,w)$ satisfies the following conditions:

1. *$P(z,w) = \sum_{m,n\geq 0} p_{m,n} z^m w^n$ is analytic*[9] *in a domain $|z| < R$ and $|w| < S$ for some $R, S > 0$.*
2. *The coefficients of P satisfy*

$$p_{m,n} \geq 0, \qquad p_{0,0} = 0, \qquad p_{0,1} \neq 1,$$

   and $p_{m,n} > 0$ for some m and for some $n \geq 2$.

3. *There exist (then necessarily unique) two numbers r, s, such that $0 < r < R$ and $0 < s < S$, satisfying the system of equations*

$$\begin{aligned} P(r,s) &= s, \\ P_w(r,s) &= 1, \end{aligned}$$

   which is called the characteristic system*, where P_w is the derivative of $P(z,w)$ with respect to w.*

Theorem 4.17 (Theorem VII.3 in [FS09], restated[10])**.** *Let $y(z)$ belong to the smooth implicit-function schema defined by $P(z,w)$, with (r,s) the positive solution of the characteristic system. Then,*

$$y(z) \bowtie r^{-1}.$$

5 The Zig-Zag product

In this section we prove both our lower and upper bounds on the Zig-Zag product (see Section 5.1 and Section 5.3, respectively). Following the lower bound proof, in Section 5.2, we deduce our universal bound which holds for all degree-c graphs. We end this section by observing the existence of "trivial" eigenvalues of the Zig-Zag product when using a

[9] Recall that a bivariate function $f(x,y)$ is analytic in a point (a,b) if and only if the functions $f_b(x) = f(x,b)$ and $f_a(y) = f(a,y)$ are analytic in the points $x = a$ and $y = b$ correspondingly, as univariate functions.

[10] The original theorem specifies much more than the asymptotic exponent, but this restatement of the theorem suffices for our needs.

consistently-labeled graph G (see Section 5.4), and analyze the Zig-Zag product of several specific graphs in Section 5.5.

5.1 Lower bound

In this subsection we prove our lower bound on the spectral expansion of a Zig-Zag product. To state our main theorems, for every graph H we define the function

$$\mathcal{Z}_H(x) = \sqrt{x^2 - \frac{x}{\mathcal{G}_{H^2}(x)}}, \tag{20}$$

where $\mathcal{G}_{H^2}$ is the Cauchy transform of H^2.

Definition 5.1 (Bad graphs for the Zig-Zag product)**.** *Let H be a vertex-transitive (d, c)-graph. We say that H is* bad *if the number of connected components in H^2 is at least $\frac{d}{2}$. Otherwise, H is called* good.

Note that our classification of graphs as either "bad" or "good" applies exclusively to vertex-transitive graphs. This classification can be made more explicit. Indeed, bad graphs are divided into two distinct categories:

1. Unions of disjoint 4-cycles, and
2. Perfect matching graphs with c_1 parallel edges together with c_2 self-loops at each vertex, for some $c_1, c_2 \geq 0$ such that $c_1 + c_2 = c$.

The Zig-Zag product of any graph G with a bad graph is degenerate in some sense, and is amenable to a straightforward analysis, which we leave to the reader.

Theorem 5.2. *Let H be a good vertex-transitive (d, c)-graph with $d \geq 3$ and $c \geq 2$. Then, for every (n, d)-graph G,*

$$\lambda(G \textcircled{z} H) \geq \min_{x > c^2} \mathcal{Z}_H(x) - o_n(1).$$

The proofs of Theorem 5.2 makes use of the trace method (see Lemma 4.5), and hence lower bounding the spectral expansion is reduced to lower bounding the number of cycles of any size in a certain infinite cover of the product. For lower bounding the number of cycles in the cover, we separate the analysis into two primary components: deducing the symbolic relation for the class of cycles in the infinite cover, resulting in a functional equation for our desired generating function; and applying analytic combinatorics, in order to determine the asymptotics of our desired combinatorial class. This is a very general method which finds itself useful here, as well as in Section 6 and in [CCMP23] in which similar objects were analyzed.

5.1.1 Symbolic analysis

Let H be vertex-transitive (d, c)-graph. Fix an arbitrary vertex v in H, and let $\widetilde{H}_v$ be the graph obtained from H^2, where the steps using v as a middle vertex are omitted; that is, if $\{u, v\}$ and $\{v, w\}$ are edges in H, then the walk $u \to v \to w$ will not induce the edge $\{u, w\}$ in $\widetilde{H}_v$. Let $\Gamma(v)$ be the set of neighbors of v in H. We use the infinite cover of graphs of the form $G \text{ⓩ} H$, and work with the graph $\mathcal{T}_d \text{ⓩ} H$. We define a truncated version of this infinite object so it will be more amenable to analyze with analytic combinatorics. Fix an arbitrary vertex $r = (v_{\mathcal{T}_d}, v_H)$ in $\mathcal{T}_d \text{ⓩ} H$. A single Zig-Zag step corresponds to a length-3 walk of the form

$$(v_{\mathcal{T}_d}, v_H) \underset{H-\mathtt{step}}{\longrightarrow} (v_{\mathcal{T}_d}, v'_H) \underset{\mathcal{T}_d-\mathtt{step}}{\longrightarrow} (u_{\mathcal{T}_d}, v''_H) \underset{H-\mathtt{step}}{\longrightarrow} (u_{\mathcal{T}_d}, u_H).$$

There are potentially c distinct values that $u_{\mathcal{T}_d}$ may attain, or less in case of parallel edges in H. Choose an initial H-step arbitrarily and truncate all the other edges in which r participates. The chosen H-step uniquely determines $u_{\mathcal{T}_d}$. Now, truncate all the edges that connect $(u_{\mathcal{T}_d}, \cdot)$ with $(v_{\mathcal{T}_d}, \cdot)$ except the ones that connect to r. We call this truncated graph X. See Figure 1 for an illustration.

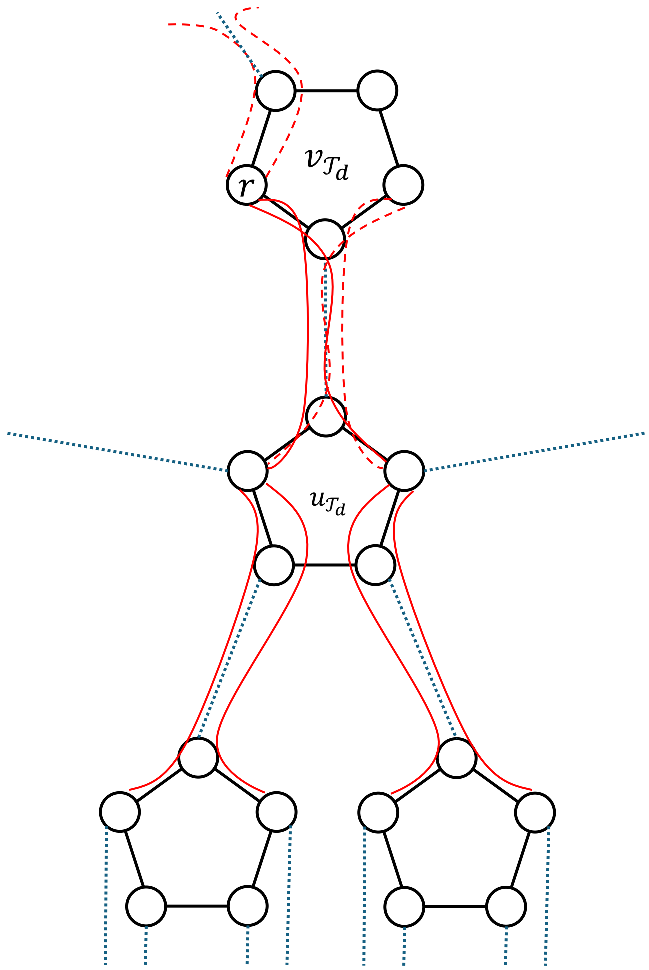

Figure 1: The Zig-Zag product graph $\mathcal{T}_5 ⓩ C_5$. The black straight solid edges are those of C_5. The blue dotted edges come from the rotation map of $\mathcal{T}_5$. The curly edges are those of the Zig-Zag product (note that they follow a Zig-Zag path: one step on C_5, followed by a step on $\mathcal{T}_5$ and then a second C_5-step). In X, the dashed curly edges are truncated.

Towards counting the cycles in X, we define three combinatorial classes:

1. $\mathcal{A}_H$: Paths from $\Gamma(v)$ to $\Gamma(v)$ in $\widetilde{H}_v$, where $\Gamma(v)$ is the multi-set of neighbors of v in H. The corresponding generating function is defined by $A_H(z)$.

2. $\mathcal{S}_{H^2}$: Cycles in H^2, originating at the vertex v, returning to v only upon completing the cycle (not including the empty cycle). The corresponding generating function is defined by $S_{H^2}(z)$.

3. $\mathcal{S}_X$: Cycles in X, originating at the vertex r, returning to r only upon completing the cycle (not including the empty cycle).The corresponding generating function is defined by $S_X(z)$.

The class of arbitrary cycles originating at the root vertex r in X is given by $\mathcal{C}_X := \mathsf{SEQ}(\mathcal{S}_X)$. It is not hard to prove, using tools from analytic combinatorics, that the generating functions $C_X(z)$ and $S_X(z)$ share their radius of convergence and hence the growth rates of their coefficients are the same, up to sub-exponential factors. For more on this subject, see *sub-criticality* in Chapter VI.9 in [FS09]. We thus focus on finding the asymptotic behavior of the coefficients of $S_X(z)$, which by definition lower bound the coefficients of $C_X(z)$. We will first state and prove some lemmata towards this end.

Claim 5.3. *For the classes $\mathcal{S}_{H^2}$ and $\mathcal{A}_H$ defined above, we have*

$$S_{H^2}(z) = 1 - \frac{1}{\frac{1}{z}\mathcal{G}_{H^2}(\frac{1}{z})},$$

and

$$A_H(z) = \frac{S_{H^2}(z)}{z} = \frac{1}{z} - \frac{1}{\mathcal{G}_{H^2}(\frac{1}{z})}. \tag{21}$$

Furthermore, the radius of convergence of $A_H(z)$ is lower bounded by $\frac{1}{c^2}$.

Proof. The first part was proven in [CCMP23], and we provide the proof here for completeness. The combinatorial class $\mathcal{C}_{H^2}$ of arbitrary cycles originated at v is related to the class $\mathcal{S}_{H^2}$ by $\mathcal{C}_{H^2} = \mathsf{SEQ}(\mathcal{S}_{H^2})$. Indeed, any nonempty cycle is a sequence of cycles returning to v exactly once, and the empty cycle is correctly captured by the SEQ construct. The relation between the generating functions of $\mathcal{C}_{H^2}$ and $\mathcal{S}_{H^2}$ is thus given by $C_{H^2}(z) = \frac{1}{1-S_{H^2}(z)}$, or equivalently, $S_{H^2}(z) = 1 - \frac{1}{C_{H^2}(z)}$. Therefore, it suffices to prove that $C_{H^2}(z) = \frac{1}{z}\mathcal{G}_{H^2}(\frac{1}{z})$. To this end, note that as H is vertex-transitive, $C_{H^2}(z)$ can also be written as

$$C_{H^2}(z) = e_v^{\mathsf{T}}(\mathbf{I} - z\mathbf{H}^2)^{-1}e_v = \frac{1}{d}\operatorname{Tr}\left(\left(\mathbf{I} - z\mathbf{H}^2\right)^{-1}\right),$$

where e_v denotes the vector satisfying $(e_v)_u = 0$ for all $u \neq v$ and $(e_v)_v = 1$. Now,

$$\mathcal{G}_{H^2}(x) = \frac{1}{d}\operatorname{Tr}\left(\left(x\mathbf{I} - \mathbf{H}^2\right)^{-1}\right) = \frac{1}{xd}\operatorname{Tr}\left(\left(\mathbf{I} - x^{-1}\mathbf{H}^2\right)^{-1}\right).$$

Substituting $x = \frac{1}{z}$, we see that $C_{H^2}(z) = \frac{1}{z}\mathcal{G}_{H^2}(\frac{1}{z})$, which completes the proof of the first part.

For the second part, Equation (21), we observe that there is a one-to-one correspondence between simple paths from v to v in H^2 and paths from $\Gamma(v)$ to $\Gamma(v)$ in $\widetilde{H}_v$, in which their sizes only differ by 1: taking a path of the former and breaking it down as a walk in H with twice as many steps, looking at the odd vertices defines a path from some $i \in \Gamma(v)$ to some $j \in \Gamma(v)$ in $\widetilde{H}_v$.

As for the radius of convergence of $A_H(z)$, the singularity of $A_H(z)$ at $z = 0$ is a removable singularity (indeed, recall from Equation (21) that $A_H(z) = \frac{S_{H^2}(z)}{z}$, and note that

$[z^0]S_{H^2}(z) = 0$ as the class $\mathcal{S}_{H^2}$ does not include the empty cycle), and hence does not affect the radius of convergence. Observing Equation (21) again, we see that the singularities of $A_H(z)$ must thus come from the singularities of $\frac{1}{\mathcal{G}_{H^2}\left(\frac{1}{z}\right)}$. The only singularities of $\mathcal{G}_{H^2}\left(\frac{1}{z}\right)$ are poles, which translate to removable singularities of $\frac{1}{\mathcal{G}_{H^2}\left(\frac{1}{z}\right)}$. The singularities of $A_H(z)$ must thus come from the zeros of $\mathcal{G}_{H^2}\left(\frac{1}{z}\right)$. By the identity

$$\mathcal{G}_{H^2}(x) = \frac{1}{d} \cdot \frac{\chi_x'(H^2)}{\chi_x(H^2)},$$

the zeros of $\mathcal{G}_{H^2}(x)$ are easily seen to be a subset of the roots of $\chi_x'(H^2)$. These roots are all real and their magnitude is bounded above by c^2, since $\chi_x'(H)$ is interlacing with $\chi_x(H^2)$. Hence, the radius of convergence of A_H is lower bounded by $\frac{1}{c^2}$, as desired. □

Lemma 5.4. *The class $\mathcal{S}_X$ satisfies the symbolic relation*

$$\mathcal{S}_X = \mathcal{Z} \times (\mathcal{A}_H \circ \mathcal{S}_X) \times \mathcal{Z},$$

where $\mathcal{Z}$ is an atomic class. In particular, $S_X(z)$ satisfies the functional equation

$$S_X(z) = z^2 \cdot A_H(S_X(z))$$

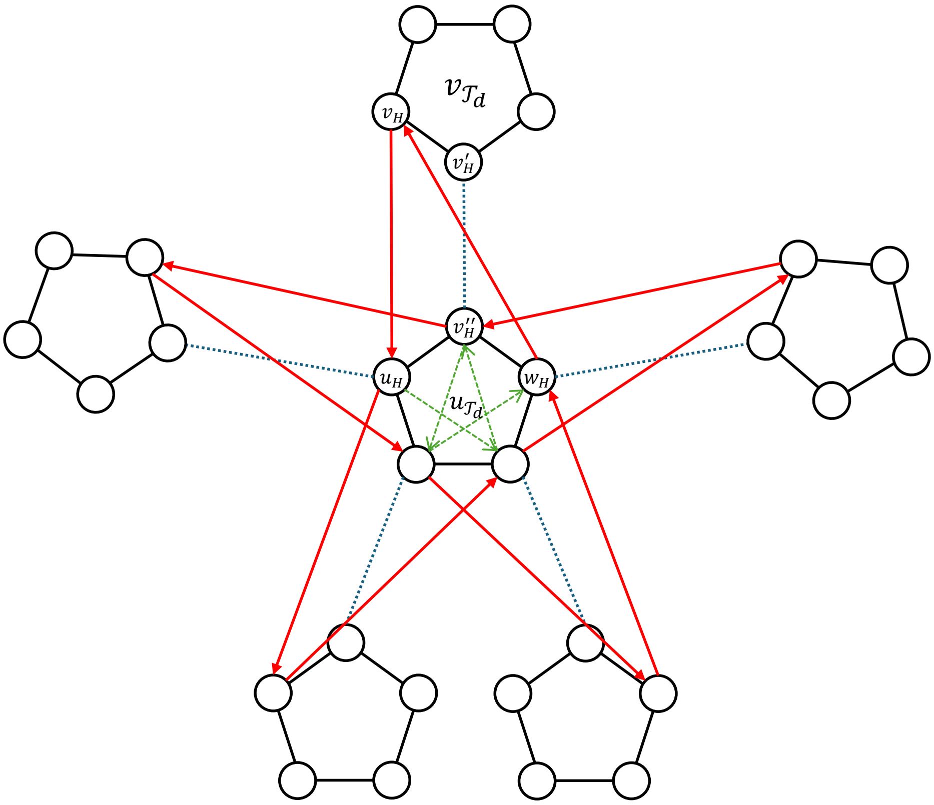

Figure 2: A cycle in the Zig-Zag product $\mathcal{T}_5 \textcircled{z} C_5$. The black solid edges are those of C_5. The blue dotted edges come from the rotation map of $\mathcal{T}_5$. The red solid arrows are valid steps in the Zig-Zag product. The green dashed arrows are steps on the cloud of $u_{\mathcal{T}_d}$, according to the transition matrix $\widetilde{H}_v$.

Proof. Consider a cycle $C \in \mathcal{S}_X$. Denote the first step in C by $(v_{\mathcal{T}_d}, v_H) \to (u_{\mathcal{T}_d}, u_H)$ and the last step by $(u_{\mathcal{T}_d}, w_H) \to (v_{\mathcal{T}_d}, v_H)$. Per our truncation, the first Zig-Zag step must be of the form

$$(v_{\mathcal{T}_d}, v_H) \underset{H-\texttt{step}}{\longrightarrow} (v_{\mathcal{T}_d}, v'_H) \underset{\mathcal{T}_d-\texttt{step}}{\longrightarrow} (u_{\mathcal{T}_d}, v''_H) \underset{H-\texttt{step}}{\longrightarrow} (u_{\mathcal{T}_d}, u_H),$$

and the last Zig-Zag step must be of the form

$$(u_{\mathcal{T}_d}, w_H) \underset{H-\texttt{step}}{\longrightarrow} (u_{\mathcal{T}_d}, v''_H) \underset{\mathcal{T}_d-\texttt{step}}{\longrightarrow} (v_{\mathcal{T}_d}, v'_H) \underset{H-\texttt{step}}{\longrightarrow} (v_{\mathcal{T}_d}, v_H).$$

for some $u_H, w_H \in \Gamma_H(v''_H)$. These two steps correspond to the two atomic $\mathcal{Z}$ elements in the symbolic relation stated in Lemma 5.4. As for the middle steps, let us look at the projection of the cycle C to visits in the cloud of $u_{\mathcal{T}_d}$, (vertices in $\{u_{\mathcal{T}_d}\} \times V_H$). We claim that the projection of C to these vertices is isomorphic to a path in $\mathcal{A}_H$. As H is vertex-transitive, paths from $\Gamma(v''_H)$ to $\Gamma(v''_H)$ in $\widetilde{H}_{v''_H}$ have one-to-one correspondence with paths from $\Gamma(v)$ to $\Gamma(v)$ in $\widetilde{H}_v$. We thus turn to show that the projection of C is a walk on $\widetilde{H}_{v''_H}$

which starts and ends in $\Gamma(v''_H)$. As $w_H, u_H \in \Gamma(v''_H)$, the above projection of the cycle C indeed starts and ends in $\Gamma(v''_H)$. As for the steps in-between, we claim that the steps in $\widetilde{H}_{v''_H}$ are exactly all the possible steps. Indeed, consider a Zig-Zag step from $(u_{\mathcal{T}_d}, \cdot)$. It must be of the form

$$(u_{\mathcal{T}_d}, x_H) \underset{H-\texttt{step}}{\longrightarrow} (u_{\mathcal{T}_d}, y_H) \underset{\mathcal{T}_d-\texttt{step}}{\longrightarrow} (t_{\mathcal{T}_d}, \cdot) \underset{H-\texttt{step}}{\longrightarrow} (t_{\mathcal{T}_d}, \cdot). \tag{22}$$

As $\mathcal{T}_d$ is a tree, in order to go back to $(u_{\mathcal{T}_d}, \cdot)$, we must perform a step of the form

$$(t_{\mathcal{T}_d}, \cdot) \underset{H-\texttt{step}}{\longrightarrow} (t_{\mathcal{T}_d}, \cdot) \underset{\mathcal{T}_d-\texttt{step}}{\longrightarrow} (u_{\mathcal{T}_d}, y_H) \underset{H-\texttt{step}}{\longrightarrow} (u_{\mathcal{T}_d}, z_H). \tag{23}$$

Projected to the cloud of $u_{\mathcal{T}_d}$, we stepped from x_H to z_H. That is, stepped on H^2. However, note that if $y_H = v''_H$ then per our truncation, Equation (22) must terminate in r. As cycles in $\mathcal{S}_X$ visit r only upon completing the cycle, such a step cannot correspond to a middle step in C, which explains the use of $\widetilde{H}_{v''_H}$ instead of H^2.

So far we showed that the cycle C starts with a step that takes us from r to the cloud of $u_{\mathcal{T}_d}$ and ends with a step that takes us back from that cloud to r. We showed that projected to this cloud, the cycle C is in a one-to-one correspondence with a path in $\mathcal{A}_H$. We now claim that each such middle step may be substituted with a path that is in one-to-one correspondence with the cycles in $\mathcal{S}_X$. To illustrate this, consider a step $x \to z$ on $\widetilde{H}_{v''_H}$ which is induced by the length-two walk on H given by $x \to y \to z$. The first H-step $x \to y$ requires us to take a detour through a specific cloud (given by the first coordinate of $\mathsf{Rot}(u_{\mathcal{T}_d}, y)$). This specific cloud is in correspondence with the unique cloud which we choose not to truncate in our truncation process. We can think of it as a determination of the first H-step in Equation (22), which ultimately determines also the $\mathcal{T}_d$-step. The second H-step $y \to z$ determines the last H-step in Equation (23), in addition to the last $\mathcal{T}_d$-step which is predetermined. This determination corresponds to the second part of our truncation where we only allow to go back to r instead of going back to other vertices in the same cloud. As a consequence, each cycle $C' \in \mathcal{S}_X$ corresponds to walks of the form $x \to y \to x$. Substituting only the last H-step by $y \to z$, we get a one-to-one correspondence between cycles in $\mathcal{S}_X$ and instantiations of the step $x \to z$ as paths in X. $\square$

In the last section we completed the symbolic derivation of the relevant classes for our analysis. We now tend to analyze $S_X(z)$. As $\mathcal{T}_d$ is bipartite and as each step in the Zig-Zag product corresponds to a single step on $\mathcal{T}_d$, the Zig-Zag product $\mathcal{T}_d ⓩ H$ is bipartite and thus all cycles in $\mathcal{S}_X$ are of even sizes. Hence, defining the function E by

$$[z^n]E(z) := [z^{2n}]S_X(z),$$

we obtain the relation $E(z^2) = S_X(z)$ and so by Lemma 5.4,

$$E(z) = z \cdot A_H(E(z)). \tag{24}$$

E is now amenable to analysis using Theorem 4.15. In order to invoke this theorem, we first need to prove that $E(z)$ belongs to the smooth inverse-function schema. In particular, we need to solve the characteristic equation induced by $E(z)$, which is given by $A_H(u) = u \cdot A'_H(u)$.

5.1.2 Solving the characteristic equation and minimizing $\mathcal{Z}_H$

In this section we will show that $\mathcal{Z}_H(x)$ must attain a minimum in the range $x > c^2$, and then relate its minimizer to the solution of the characteristic equation $A_H(u) = u \cdot A'_H(u)$. To this end, it suffices to prove that the term inside the square root in the definition of $\mathcal{Z}_H(x)$,

$$g(x) := x^2 - \frac{x}{\mathcal{G}_{H^2}(x)},$$

has a positive minimum in this range. In Lemma 5.8 we bridge between the minimum of $\mathcal{Z}_H$ and the solution to the characteristic equation defined by Equation (24), $A_H(u) = u \cdot A'_H(u)$.

Claim 5.5. *The function $g(x)$ has a positive minimum in the range $x > c^2$.*

Proof. As $g(x)$ and its derivative are continuous in the range $x > c^2$, it suffices to prove that near c^2, $g'(x) < 0$ and that $g(x) \to \infty$ as $x \to \infty$. In both proofs, we will need to investigate the Cauchy transform in order to analyze the terms. Denote by $\lambda_1 \geq \cdots \geq \lambda_d$ the eigenvalues of the adjacency matrix of H. The eigenvalues of H^2 are thus $\{\lambda_i^2\}_{i=1}^d$. Writing $\mathcal{G}_{H^2}$ explicitly, we have

$$\mathcal{G}_{H^2}(x) = \frac{1}{d} \sum_{i=1}^{d} \frac{1}{x - \lambda_i^2}.$$

Claim 5.6.

$$\lim_{x \to c^2} g'(x) < 0.$$

Proof.

$$g'(x) = 2x + \frac{x\mathcal{G}'_{H^2}(x)}{\mathcal{G}^2_{H^2}(x)} - \frac{1}{\mathcal{G}_{H^2}(x)}. \tag{25}$$

Taking the limit of $g'(x)$ as x approaches c^2, $2x$ approaches $2c^2$ and $\frac{1}{\mathcal{G}_{H^2}(x)}$ approaches 0. Now,

$$\frac{\mathcal{G}'_{H^2}(x)}{\mathcal{G}^2_{H^2}(x)} = -\frac{\frac{1}{d}\sum_{i=1}^{d}\left(\frac{1}{x-\lambda_i^2}\right)^2}{\left(\frac{1}{d}\sum_{i=1}^{d}\frac{1}{x-\lambda_i^2}\right)^2} = -d\frac{\sum_{i=1}^{d}\left(\frac{x-c^2}{x-\lambda_i^2}\right)^2}{\left(\sum_{i=1}^{d}\frac{x-c^2}{x-\lambda_i^2}\right)^2}. \tag{26}$$

Denote by m_1 the multiplicity of the eigenvalue c^2 in H^2. As x approaches c^2, this ratio approaches

$$-\frac{\frac{m_1}{d}}{\frac{m_1^2}{d^2}} = -\frac{d}{m_1}.$$

It is well known that m_1 equals to the number of connected components in H^2. As we assumed H is good (in the sense of Definition 5.1), $m_1 < \frac{d}{2}$. Plugging this back, we can see that as x approaches c^2, $g'(x)$ approaches $2c^2 - \frac{d}{m_1} \cdot c^2 < 0$, proving Claim 5.6. ☐

Claim 5.7. *In the range $x > c^2$, $g(x)$ is strictly positive and, satisfies*

$$\lim_{x \to \infty} g(x) = \infty.$$

Proof. For all $x > c^2 = \lambda_1^2$, we have that $x - \lambda_i^2$ is positive for all i and hence $\frac{1}{x-\lambda_i^2} \geq \frac{1}{x}$. Consequently, by applying this bound for all $i \neq 1$, we get

$$\mathcal{G}_{H^2}(x) = \frac{1}{d} \sum_{i=1}^{d} \frac{1}{x - \lambda_i^2} \geq \frac{d-1}{d} \cdot \frac{1}{x} + \frac{1}{d} \cdot \frac{1}{x - c^2}.$$

So for all $x > c^2$,

$$g(x) = x^2 \left(1 - \frac{1}{x\mathcal{G}_{H^2}(x)}\right) \geq x^2 \left(1 - \frac{1}{\frac{d-1}{d} + \frac{1}{d} \cdot \frac{x}{x-c^2}}\right) = \frac{c^2 x^2}{dx - c^2(d-1)}.$$

It is now easy to see that in the domain $x > c^2$, the last terms is positive, and diverges to ∞ as $x \to \infty$, which concludes the proof. ☐

Claim 5.6 and Claim 5.7 and the trivial continuity of $g'(x)$ in the domain $x > c^2$ imply the existence of a point $x_0 > c^2$ which minimizes $g(x)$, and that $g(x_0) > 0$. In particular, $g'(x_0) = 0$. ☐

Lemma 5.8. *For $x_0 > c^2$, the following conditions are equivalent:*

- $g'(x_0) = 0$, *and*
- $\frac{1}{x_0}$ *solves the characteristic equation $A_H(u) = u \cdot A'_H(u)$.*

Proof. Let us write A_H in terms of $\mathcal{G}_{H^2}$. We have that

$$A_H(u) = \frac{1}{u} - \frac{1}{\mathcal{G}_{H^2}\left(\frac{1}{u}\right)},$$

$$A'_H(u) = -\frac{1}{u^2} - \frac{\mathcal{G}'_{H^2}\left(\frac{1}{u}\right)}{\mathcal{G}^2_{H^2}\left(\frac{1}{u}\right) \cdot u^2}.$$

Substituting variables $u = \frac{1}{x}$ and writing the characteristic equation in terms of x instead of u, we get the equation

$$0 = A_H\left(\frac{1}{x}\right) - \frac{1}{x}A'_H\left(\frac{1}{x}\right) = 2x - \frac{1}{\mathcal{G}_{H^2}(x)} + x\frac{\mathcal{G}'_{H^2}(x)}{\mathcal{G}^2_{H^2}(x)} = g'(x),$$

which completes the proof of the lemma. $\square$

Corollary 5.9. *There exists a solution $\tau \in \left(0, \frac{1}{c^2}\right)$ to the characteristic equation $A_H(u) = u \cdot A'_H(u)$, within the radius of convergence of $A_H(u)$.*

Proof. The existence of the solution to the characteristic equation within the domain $\left(0, \frac{1}{c^2}\right)$ is implied by Claim 5.5 and Lemma 5.8. By Claim 5.3, since the radius of convergence of $A_H(u)$ is lower bounded by $\frac{1}{c^2}$, the solution lies within its radius of convergence of $A_H(u)$, as desired. $\square$

Corollary 5.10. *There is a unique point $x_0 \in (c^2, \infty)$ such that $\mathcal{Z}'_H(x_0) = 0$. In x_0, $\mathcal{Z}_H(x)$ attains its global minimum in that domain.*

Proof. By Theorem 4.15, the solution $\tau \in \left(0, \frac{1}{c^2}\right)$ to the characteristic equation $A_H(u) = u \cdot A'_H(u)$ is unique. Lemma 5.8 then concludes the proof. $\square$

5.1.3 Coefficients extraction

In the previous sections we translated the problem of approximating the number of cycles of each length in X to estimating the coefficients in the series expansion of the function $E(z)$. We reduced the problem of solving the characteristic equation to the problem of minimizing $\mathcal{Z}_H(x)$. We are now ready to conclude Theorem 5.2.

Claim 5.11. *The function $E(z)$ belongs to the smooth inverse-function schema.*

Proof. As long as $\widetilde{H}_v$ has an edge $e \in \Gamma(v) \times V_H$, paths in $\mathcal{A}_H$ of every even length may be produced by walking back and forth along e. In particular, all the even coefficients are non-zero and hence $A_H(u) \neq A_0 + A_1 u$. It is easy to verify that the assumption that H is good implies that this is indeed the case. Moreover, $A_H(0) \geq c > 0$. Indeed, there are c elements in $\Gamma(v)$, each one contributes an empty path from $\Gamma(v)$ to $\Gamma(v)$. In case of parallel edges in H, $A_H(0)$ might be even larger. In any case, $A_H(0) \neq 0$. Finally, by the definition of $A_H(z)$ as a generating function of a combinatorial class that counts paths, its coefficients must be non-negative. Thus, the conditions in Equation (19) are satisfied.

The solution to the characteristic equation $A_H(u) = u \cdot A'_H(u)$ within the radius of convergence of $A_H(u)$ was already established in Corollary 5.9, concluding that $E(z)$ belongs to the smooth inverse-function schema. $\square$

We are now ready to prove Theorem 5.2.

Proof. By Claim 5.11, we get that $E(z)$ belongs to the smooth inverse-function schema. By *Lemma* 5.8, the solution τ to the characteristic equation $A_H(u) = u \cdot A'_H(u)$ is such that $x_0 = \frac{1}{\tau}$ is the unique minimizer of $\mathcal{Z}_H(x)$ in the domain (c^2, ∞). Now, we have that

$$A'_H(\tau) = \frac{A_H(\tau)}{\tau} = \frac{1}{\tau^2} - \frac{1}{\tau \mathcal{G}_{H^2}\left(\frac{1}{\tau}\right)} = x_0^2 - \frac{x_0}{\mathcal{G}_{H^2}(x_0)}.$$

Hence, by Corollary 5.10,

$$\sqrt{A'_H(\tau)} = \sqrt{x_0^2 - \frac{x_0}{\mathcal{G}_{H^2}(x_0)}} = \min_{x>c^2} \mathcal{Z}_H(x).$$

Invoking Theorem 4.15, we get that $E \bowtie A'_H(\tau)$. By the definition of $E(z)$, we get that

$$S_X \bowtie \sqrt{A'_H(\tau)} = \min_{x>c^2} \mathcal{Z}_H(x).$$

Note that $[z^n]S_X(z)$ is a lower bound on the number of cycles in $\mathcal{T}_d \circledz H$, both because the truncation reduces the number of cycles and because $\mathcal{S}_X$ only counts cycles that visit the root vertex r only upon completing a cycle. Invoking Lemma 4.5, we conclude that for every graph G on n vertices,

$$\lambda_2(G \circledz H) \geq (1 - o_n(1)) \min_{x>c^2} \mathcal{Z}_H(x) = \min_{x>c^2} \mathcal{Z}_H(x) - o_n(1).$$

$\square$

5.2 Universal Zig-Zag lower bound depending on degree only

Theorem 5.2 gives a precise analytical viewpoint of how the spectrum of H^2, and equivalently its cycle structure, influences the spectrum of $G \circledz H$. They require, however, full understanding of H's structure and solving a separate minimization problem for every H of interest. The following theorem captures the best possible expansion one can expect from a Zig-Zag product, and depends *only* on the degree c of the small graph H. It also fully captures the asymptotic potential and limitation of the Zig-Zag product. We further analyze it for the specific case of $c = 2$ and in particular cycle graphs in Section 5.5.1.

Theorem 5.12. *Let H be a good vertex-transitive (d, c)-graph with $d \geq 3$ and $c \geq 3$. Then, for every (n, d)-graph G,*

$$\lambda(G \circledz H) \geq \frac{c^2}{\sqrt{c-1}} - o_n(1). \tag{27}$$

For the particular case of $c = 2$,

$$\lambda(G ⓩ H) \geq \frac{3\sqrt{3}}{\sqrt{2}} - o_n(1) \approx 3.674.$$

A clear consequence of Theorem 5.12 is that a Zig-Zag product of two graphs is always at least quadratically far from being Ramanujan. It is worth noting that the parameter d, the degree of the large graph and size of the small one, does not play a role in Equation (27). This is for the reason that, as the Alon-Boppana bound can be achieved by counting cycles in the infinite d-regular tree, we think similarly of H's size as being infinity. Hence we only expect graphs to achieve this bound in the large d limit, and we leave open the question of a universal lower bound depending on both c and d (but not on the full spectrum of H).

For proving Theorem 5.12, we need two additional definitions: the *Kesten-McKay* distribution, which represents the spectrum of the inifinite d-regular tree, and a more general definition of the Cauchy transform.

Definition 5.13. *Let a be a real number, and let μ be a probability distribution over $[-a, a] \subset \mathbb{R}$. The* Cauchy transform *of μ is the function $\mathcal{G}_\mu : (a, \infty) \to \mathbb{R}$ defined by:*

$$\mathcal{G}_\mu(x) = \int_{\mathbb{R}} \frac{1}{x - t} \mu(t) dt.$$

The definition shown in Section 4.2.1 is a special case of the above, when μ is considered to be the (discrete) uniform distribution over the roots of the polynomial $p(x)$. It is also known that for every $x > a$,

$$\mathcal{G}_\mu(x) = \sum_{r=0}^{\infty} \frac{m_r(\mu)}{x^{r+1}}, \tag{28}$$

$m_r(\mu)$ being the r-th moment of μ.

Definition 5.14. *The Kesten-McKay distribution with parameter c is given by the probability density function*

$$\mu_{\mathsf{km}}(t) = \begin{cases} \dfrac{c\sqrt{4(c-1) - t^2}}{2\pi(c^2 - t^2)}, & \text{for } |t| \leq 2\sqrt{c-1}; \\ 0, & \text{otherwise.} \end{cases} \tag{29}$$

This distribution is known (see, e.g., [McK81]) to represent the spectrum of the infinite c-regular tree $\mathcal{T}_c$. In particular, its moments represent the number of closed walks of every length originating at the root of $\mathcal{T}_c$, and so for positive x we have by Equation (28) that

$$\mathcal{G}_H(x) \geq \mathcal{G}_{\mathsf{km}}(x) \quad \text{and} \quad \mathcal{G}_{H^2}(x) \geq \mathcal{G}_{\mathsf{km}^2}(x), \tag{30}$$

where by km^2 we denote the distribution of a random variable X^2, where X is distributed according to μ_{km}. It can easily be derived (see, e.g. [CM23] Section 3.3) that

$$\mathcal{G}_{\mathsf{km}^2}(x) = \frac{c\sqrt{x-4(c-1)} - \sqrt{x}(c-2)}{2\sqrt{x}(x-c^2)}. \tag{31}$$

Proof of Theorem 5.12. Let

$$\Psi(x,y) = \sqrt{x^2 - \frac{x}{y}}.$$

We note that for $x > 0$, $\Psi(x,y)$ is monotone increasing in y and that $\mathcal{Z}_H(x) = \Psi\left(x, \mathcal{G}_{H^2}(x)\right)$. We denote

$$\mathcal{Z}_{\mathcal{T}_c}(x) = \Psi\left(x, \mathcal{G}_{\mathsf{km}^2}(x)\right).$$

By Equation (30) we get that $\mathcal{Z}_H(x) \geq \mathcal{Z}_{\mathcal{T}_c}(x)$ for all $x > c^2$. By a direct calculation,

$$\mathcal{Z}_{\mathcal{T}_c}(x) = \sqrt{\frac{c}{2c-2}} \cdot \sqrt{x^2 - x^{3/2}\sqrt{4-4c+x}}. \tag{32}$$

The global minimal value of $\mathcal{Z}_{\mathcal{T}_c}(x)$ is $\mathcal{Z}_{\mathcal{T}_c}(x_0) = \frac{3\sqrt{3}}{2}\sqrt{c(c-1)}$, achieved at $x_0 = \frac{9(c-1)}{2}$. Therefore,

$$\lambda(G \textcircled{z} H) \geq \frac{3\sqrt{3}}{2}\sqrt{c(c-1)} - o_n(1). \tag{33}$$

Note that, as $\frac{9(c-1)}{2} < c^2$ for $c \geq 4$ (with equality holding for $c = 3$), x_0 is not a valid assignment for the minimization problem $\min_{x>c^2} \mathcal{Z}_H(x)$.

It can be easily verified that $\mathcal{Z}_{\mathcal{T}_c}(x)$ is increasing for $x > x_0$, and therefore the optimal bound for the minimization problem is achieved at $x = c^2$, yielding the lower bound $\mathcal{Z}_{\mathcal{T}_c}(c^2) = \frac{c^2}{\sqrt{c-1}}$. $\square$

An example of the above, comparing the universal Zig-Zag lower bound with $c = 3$ to the lower bound on Zig-Zag with the Petersen graph, is shown in Figure 3.

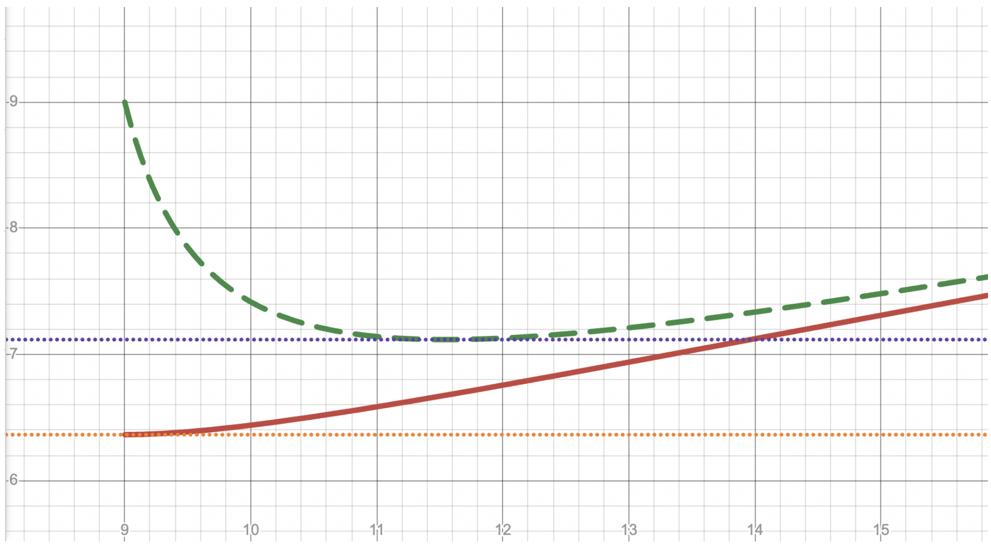

Figure 3: The green dashed line represents $\mathcal{Z}_{\mathrm{Pet}}(x)$ as defined in Section 2. The red line is $\mathcal{Z}_{T_3}(x)$. The dotted purple line is the lower bound for $\lambda(G ⓩ \mathrm{Pet})$, and the orange one is the universal lower bound for $\lambda(G ⓩ H)$ when H is of degree 3.

It is worth noting that the bound shown in Equation (33) can be proved using elementary combinatorial means, by carefully counting the cycles of the Zig-Zag of two infinite trees. We add this self contained proof of this weaker version of Theorem 5.12 in Appendix B.

5.3 Existence of graphs matching the Zig-Zag lower bound

The goal of this section is to prove the tightness of Theorem 5.2, which we do, as discussed in Section 3.3, by defining a distribution over graphs and applying finite free probability. As we have seen in Definition 4.4, the adjacency matrix $\mathbf{Z}$ of $G ⓩ H$ can be written as $\mathbf{Z} = \boldsymbol{\mathcal{H}}\dot{\mathbf{G}}\boldsymbol{\mathcal{H}}$. By Definition 4.2, drawing G according to the configuration model is equivalent to picking $\dot{\mathbf{G}}$ as a random matching. We can now define the following distribution on graphs: draw G according to the configuration model, and return $G ⓩ H$. The corresponding (random) adjacency matrix is given by

$$\mathbf{Z}_{\mathbf{P}} = \boldsymbol{\mathcal{H}}\mathbf{P}\mathbf{M}\mathbf{P}^{\intercal}\boldsymbol{\mathcal{H}},$$

where $\mathbf{M}$ is some fixed matching of dimension nd and $\mathbf{P}$ is a uniformly random permutation matrix of the same dimension. We are now ready to state the main theorem of this section.

Theorem 5.15. *For every good (d, c)-graph H and every integer $n \geq 1$, there exists an (n, d)-graph G such that*

$$\lambda_2(G ⓩ H) \leq \min_{x > c^2} \mathcal{Z}_H(x),$$

where $\mathcal{Z}_H(x)$ is as defined in Equation (20).

Theorem 5.15, together with Theorem 5.2, establishes that both results are optimal with respect to the second largest eigenvalue (Theorem 5.2 bounds negative eigenvalues as well). Note that Theorem 5.15 does not require the graph H to be vertex-transitive.

In order to prove Theorem 5.15, we follow the steps of [MSS18] and its followups [CM23] and [CCMP23], and first bound the roots of the *expected characteristic polynomial*:

$$\mathop{\mathbf{E}}_{\mathbf{P}} \chi_x\left(\mathbf{Z}_{\mathbf{P}}\right) = \mathop{\mathbf{E}}_{\mathbf{P}} \chi_x\left(\boldsymbol{\mathcal{H}}\mathbf{P}\mathbf{M}\mathbf{P}^{\mathsf{T}}\boldsymbol{\mathcal{H}}\right)$$

using Lemma 4.8 and Lemma 4.11, and then invoke Lemma 4.9 to claim that there exists a *specific graph* G having the same bound on its roots.

Proof of Theorem 5.15. Since characteristic polynomials are invariant under cyclic rotations, we have that $\chi_x\left(\boldsymbol{\mathcal{H}}\mathbf{P}\mathbf{M}\mathbf{P}^{\mathsf{T}}\boldsymbol{\mathcal{H}}\right) = \chi_x\left(\boldsymbol{\mathcal{H}}^2\mathbf{P}\mathbf{M}\mathbf{P}^{\mathsf{T}}\right)$. Thus,

$$\mathop{\mathbf{E}}_{\mathbf{P}} \chi_x\left(\mathbf{Z}_{\mathbf{P}}\right) = \mathop{\mathbf{E}}_{\mathbf{P}} \chi_x\left(\boldsymbol{\mathcal{H}}^2\mathbf{P}\mathbf{M}\mathbf{P}^{\mathsf{T}}\right) = \left(x - c^2\right)\left(p_{\boldsymbol{\mathcal{H}}^2} \boxtimes p_{\mathbf{M}}\right)(x), \tag{34}$$

where the last equality follows by Lemma 4.8. Since c^2 is the largest eigenvalue of $G \circledz H$, bounding the rest of the spectrum reduces to the task of bounding the roots of $p_{\boldsymbol{\mathcal{H}}^2} \boxtimes p_{\mathbf{M}}$.

Recall that for every $y > 0$, $\mathcal{N}_{p_{\boldsymbol{\mathcal{H}}^2} \boxtimes p_{\mathbf{M}}}(y)$ provides an upper bound on the largest root of $p_{\boldsymbol{\mathcal{H}}^2} \boxtimes p_{\mathbf{M}}$. By Lemma 4.11, we know that for every $y > 0$,

$$\mathcal{N}_{p_{\boldsymbol{\mathcal{H}}^2} \boxtimes p_{\mathbf{M}}}(y) \leq \frac{y}{y+1} \cdot \mathcal{N}_{p_{\boldsymbol{\mathcal{H}}^2}}(y) \cdot \mathcal{N}_{p_{\mathbf{M}}}(y) \leq \frac{y}{y+1} \cdot \mathcal{N}_{\boldsymbol{\mathcal{H}}^2}(y) \cdot \mathcal{N}_{\mathbf{M}}(y), \tag{35}$$

where for the last inequality we applied Claim 4.12. By a straightforward calculation, $\mathcal{M}_{\mathbf{M}}(x) = \frac{1}{x^2-1}$, and hence $\mathcal{N}_{\mathbf{M}}(y) = \sqrt{\frac{y+1}{y}}$, so Equation (35) can be rewritten as

$$\mathcal{N}_{p_{\boldsymbol{\mathcal{H}}^2} \boxtimes p_{\mathbf{M}}}(y) \leq \sqrt{\frac{y}{y+1}} \cdot \mathcal{N}_{\boldsymbol{\mathcal{H}}^2}(y). \tag{36}$$

We know that every value of the above is an upper bound on the roots of $p_{\boldsymbol{\mathcal{H}}^2} \boxtimes p_{\mathbf{M}}$, and therefore the best we can do is to minimize this function over $y > 0$. As $\mathcal{N}_{\boldsymbol{\mathcal{H}}^2}$ is the inverse under composition of $\mathcal{M}_{\boldsymbol{\mathcal{H}}^2}$, instead of minimizing over y, we will minimize over $x > c^2$, where

$$y = \mathcal{M}_{\boldsymbol{\mathcal{H}}^2}(x) = x\mathcal{G}_{\boldsymbol{\mathcal{H}}^2}(x) - 1,$$

and we also observe that $\mathcal{G}_{\boldsymbol{\mathcal{H}}^2}(x) = \mathcal{G}_{H^2}(x)$. Overall we have that

$$\alpha_2\left(\mathop{\mathbf{E}}_{\mathbf{P}} \chi_x\left(\mathbf{Z}_{\mathbf{P}}\right)\right) \leq \min_{x > c^2}\left(\sqrt{x^2 - \frac{x}{\mathcal{G}_{H^2}(x)}}\right), \tag{37}$$

where we recall that $\alpha_k(p)$ is the k-th largest root of the polynomial $p(x)$.

For proving the existence of a graph achieving the same bound as the expected polynomial, we apply Lemma 4.9 with $k = 2$, as we see by Equation (34) that $\mathbf{E_P}\, \chi_x(\mathbf{Z_P})$ aligns with the RHS of Equation (16). $\square$

5.4 Trivial eigenvalues for consistently labeled G

The Zig-Zag construction, as described in Definition 4.4, involves not only the graphs G and H, but also the edge labeling, encoded in the rotation map of G. At first sight, it is unclear whether the labeling can influence the expansion properties of the product. On an intuitive level though, a more structured labeling yields more structured Zig-Zag products, which potentially harms expansion as we associate an expander with a pseudorandom graph. We formalize this idea in this section.

Theorem 5.2 is proven tight in Section 5.3. In addition, simulations suggest that the bound represents the typical behavior of $G \textcircled{z} H$ for a random graph G drawn according to the configuration model. However, for even n, there is a different common way to draw random (n, d)-graphs, which produces consistently labeled graphs: picking d independent random perfect matchings, and union them (each perfect matching is then associated with a unique label). In the claim below we prove that, for any consistently labeled graph G, the d eigenvalues of the graph H^2 are a subset of the nd eigenvalues of $G \textcircled{z} H$, counted with multiplicities. We follow by discussing the cases in which these *trivial* eigenvalues influence the expansion properties.

Claim 5.16. *Let $G = (V_G, E_G)$ be a consistently labelled (n, d)-graph, and let α be an eigenvalue of H with corresponding eigenvector y. Then, α^2 is an eigenvalue of $G \textcircled{z} H$ with corresponding eigenvector $\mathbf{1}_n \otimes y$.*

Proof. Let (v, i) be a vertex in $G \textcircled{z} H$. By the definition of the Zig-Zag product, for every $j \in \Gamma_H(i)$ and for every $k \in \Gamma_H(j)$, (v, i) is connected to (u, k) in $G \textcircled{z} H$, for some u in V_G. Therefore, if we only look at the second coordinate, the edges in $G \textcircled{z} H$ are equivalent to those of H^2. Since in the vector $\mathbf{1}_n \otimes y$ the values are determined by the second coordinate, we get that the eigenvalues translate to $G \textcircled{z} H$ as well. $\square$

When is Claim 5.16 making an impact on expansion properties of the graph? Had G been consistently labeled, this immediately tells us the following: if H is bipartite, c^2 is an eigenvalue of $G \textcircled{z} H$, meaning that the graph is disconnected (which is generally not the case if G is not consistently labeled). More generally, had H^2 had an eigenvalue close to c^2, the largest trivial eigenvalue may be larger than the lower bound of Theorem 5.2. This is seen in the case that H is a cycle of odd length, where the trivial eigenvalue becomes larger than our bound from Theorem 5.2 for $d \geq 11$. We elaborate on the specific case of cycle graphs, which also gained attention in [RVW00], in Section 5.5.1.

5.5 Zig-Zag with specific graphs

We examplify the usage of Theorem 5.2 and Theorem 5.12 by considering two concrete examples, comparing our results to the classical RVW analysis and to the Alon-Boppana bound.

5.5.1 Zig-Zag with cycles

An application of the Zig-Zag product, shown in the original paper [RVW00], is a way to transform any expander G of arbitrary degree d to a degree 4 expander, by zig-zagging with the cycle graph of size d, which we denote C_d.

Lemma 5.17 (Corollary 3.4 in [RVW00])**.** *Let d be odd and G be an (n, d)-graph with $\omega(G) < 1$. Then, $G Ⓩ C_d$ is an $(nd, 4)$-graph with $\omega(G Ⓩ C_d) < 1$.*

The range of possible values for $\omega(G Ⓩ C_d)$ remains, however, unspecified by Lemma 5.17. Here we prove lower bounds for it, which depend both on d and on whether or not G is consistently labeled.

Corollary 5.18. *For consistently labelled graph G, $\omega(G Ⓩ C_d) \geq 1 - \Theta(\frac{1}{d^2})$.*

Proof. It is well known that the normalized eigenvalues of C_d are $\cos\left(\frac{2\pi k}{d}\right)$ for $k \in \{0, \ldots, d-1\}$ (see, e.g., [HLW06]). Hence, the largest eigenvalue of C_d, denoted α, satisfies $\alpha = 1 - \Theta\left(\frac{1}{d^2}\right)$. The normalized largest eigenvalue of C_d^2 is therefore $\alpha^2 = 1 - \Theta\left(\frac{1}{d^2}\right)$. By Claim 5.16, α^2 is also an eigenvalue for $G Ⓩ C_d$. □

Corollary 5.18 tells us in particular that although Lemma 5.17 indicates that $G Ⓩ C_d$ is indeed an expander for constant d, for consistently labeled graphs the expansion vanishes with the degree. The more general case is the one where G is not necessarily consistently labelled, which follows immediately from Theorem 5.12.

Corollary 5.19. *For any (n, d)-graph G,*

$$\omega(G Ⓩ C_d) \geq \frac{3\sqrt{6}}{8} - o_n(1) \approx 0.91856. \tag{38}$$

Note that the value shown in Equation (38) is tight by our general existence result discussed in Section 5.3, and is larger than the Alon-Boppana bound of $\frac{2\sqrt{3}}{4} \approx 0.866$.

5.5.2 Zig-Zag with the Petersen graph

The *Petersen graph*, denoted Pet, is a 3-regular graph on 10 vertices. Its characteristic polynomial is $(x-3)(x+2)^4(x-1)^5$, and therefore taking $H = \text{Pet}$ gives

$$\chi_x(\text{Pet}^2) = (x-9)(x-4)^4(x-1)^5, \tag{39}$$

and $\omega(\mathrm{Pet}) = \frac{2}{3}$.

For any $(n, 10)$-graph G, the Zig-Zag product $G \textcircled{z} \mathrm{Pet}$ is a $(10n, 9)$-graph. By the Alon-Boppana bound,

$$2\sqrt{8} - o_n(1) \leq \lambda(G \textcircled{z} \mathrm{Pet}) \leq 9,$$

and we note that $2\sqrt{8} \approx 5.66$. We wish to know what the RVW bound gives us in comparison to our bound. At this point we look at the exact result from [RVW00] which says that $\omega(G \textcircled{z} H) \leq f\left(\omega(G), \omega(H)\right)$, where

$$f(a, b) = \frac{1}{2}\left((1 - b^2)a + \sqrt{(1 - b^2)^2 a^2 + 4b^2}\right).$$

Assuming that G is Ramanujan, the RVW bound can be shown to yield

$$\lambda(G \textcircled{z} \mathrm{Pet}) \leq \frac{3}{2}\left(1 + \sqrt{17}\right) \approx 7.685.$$

In the notations of Theorem 2.1, using Equation (39), one gets

$$\mathcal{Z}_{\mathrm{Pet}}(x) = \sqrt{3}x \cdot \sqrt{\frac{12 - 9x + x^2}{x(22 - 11x + x^2)}}.$$

The minimal value of $\mathcal{Z}_{\mathrm{Pet}}(x)$ gives the bound

$$\lambda(G \textcircled{z} \mathrm{Pet}) \geq \frac{3}{\sqrt{\frac{11}{33 + \sqrt[3]{11\left(275 - 4\sqrt{11}\right)} + \sqrt[3]{11\left(275 + 4\sqrt{11}\right)}}}} - o_n(1) \approx 7.1176.$$

Our universal bound of Theorem 2.2, which only takes into account the degree and not on the full information about Pet, tells us in this instance that

$$\lambda(G \textcircled{z} \mathrm{Pet}) \geq \frac{3^2}{\sqrt{3 - 1}} - o_n(1) \approx 6.364.$$

6 The replacement product

For analyzing the replacement product, we applied the previously used tools: analytic combinatorics and free probability. The proof of the upper bound follows identical steps to those of our proof of the upper bound for the Zig-Zag product from Section 5.3. The lower bound, on the other hand, turns out to be much more involved. It was our analysis of the upper bound, and the belief that the two should coincide, that assisted us in proving the lower bound.

Let H be a vertex-transitive (d,c) graph. Pivotal to our proofs is the function $\mathcal{R}_H : (c,\infty) \to \mathbb{R}$ which pops up in the proof of the upper bound and is extensively used in the proof of the lower, given by

$$\mathcal{R}_H(x) = x + \frac{\sqrt{1+4\mathcal{G}_H(x)^2}-1}{2\mathcal{G}_H(x)}, \tag{40}$$

where H is a (d,c)-graph and $\mathcal{G}_H$ is the Cauchy transform of H. In particular, we will seek to minimize this function. The proof of the existence and the uniqueness of the minimum of $\mathcal{R}_H(x)$ appears in Section 6.2.2, under some minimal conditions which we assume on H.

6.1 Construction

We follow the steps of Section 5.3 in order to establish the existence of an optimal family of graphs, each of which is obtained by taking the replacement product with a given (d,c)-graph H. The difference here is the order in which the upper and lower bounds are shown, so in this case the lower bound, proven in Section 6.2, will prove to be optimal. For an overview of the methodology behind the construction and proof, see Section 3.3.

Theorem 6.1. *For every (d,c)-graph H and every integer $n \geq 1$, there exists an (n,d)-graph G such that*

$$\lambda_2(G \circledR H) \leq \min_{x>c} \mathcal{R}_H(x).$$

As discussed in Definition 4.3, the adjacency matrix $\mathbf{R}$ of $G \circledR H$ can be written as $\mathbf{R} = \boldsymbol{\mathcal{H}} + \dot{\mathbf{G}}$. The rest of the proof proceeds exactly like Section 5.3. The distribution is defined as follows: draw an (n,d)-graph G according to the configuration model, and return $G \circledR H$. The corresponding (random) adjacency matrix is given by

$$\mathbf{R}_{\mathbf{P}} = \boldsymbol{\mathcal{H}} + \mathbf{P}\mathbf{M}\mathbf{P}^{\mathsf{T}},$$

where $\mathbf{M}$ is some fixed perfect matching of dimension nd, and $\mathbf{P}$ is a uniformly random permutation matrix of the same dimension.

Proof of Theorem 6.1. We start by looking at the *expected characteristic polynomial*

$$\mathop{\mathbf{E}}_{\mathbf{P}} \chi_x(\mathbf{R}_{\mathbf{P}}) = \mathop{\mathbf{E}}_{\mathbf{P}} \chi_x\left(\boldsymbol{\mathcal{H}} + \mathbf{P}\mathbf{M}\mathbf{P}^{\mathsf{T}}\right) = (x-(c+1))\left(p_{\boldsymbol{\mathcal{H}}} \boxplus p_{\mathbf{M}}\right)(x),$$

where the last equality follows by Lemma 4.8. Since $(c+1)$ is the largest eigenvalue of $G \circledR H$, bounding the rest of the spectrum reduces to the task of bounding the roots of $p_{\boldsymbol{\mathcal{H}}} \boxplus p_{\mathbf{M}}$.

We recall that for every $y > 0$, $\mathcal{K}_{p_{\boldsymbol{\mathcal{H}}} \boxplus p_{\mathbf{M}}}(y)$ provides an upper bound on the largest root of $p_{\mathcal{H}} \boxplus p_{\mathbf{M}}$. By Lemma 4.10, we know that for every $y > 0$,

$$\mathcal{K}_{p_{\mathcal{H}} \boxplus p_{\mathbf{M}}}(y) \leq \mathcal{K}_{p_{\mathcal{H}}}(y) + \mathcal{K}_{p_{\mathbf{M}}}(y) - \frac{1}{y} \leq \mathcal{K}_{\mathcal{H}}(y) + \mathcal{K}_{\mathbf{M}}(y) - \frac{1}{y}, \tag{41}$$

where the last inequality follows by Claim 4.12. It is easily verified that

$$\mathcal{K}_{\mathbf{M}}(y) = \frac{1 + \sqrt{1 + 4y^2}}{2y},$$

and so overall we get

$$\mathcal{K}_{p_{\mathcal{H}} \boxplus p_{\mathbf{M}}}(y) \leq \mathcal{K}_{\mathcal{H}}(y) + \frac{\sqrt{1 + 4y^2} - 1}{2y}. \tag{42}$$

By definition of the $\mathcal{K}$ transform, for every $y > 0$ the value of $\mathcal{K}_{p_{\mathcal{H}} \boxplus p_{\mathbf{M}}}(y)$ is an upper bound on $\alpha_2\left(\mathbf{E}_{\mathbf{P}}\, \chi_x\left(\mathbf{R}_{\mathbf{P}}\right)\right)$, and therefore we should aim to minimize the RHS of Equation (42) over $y > 0$.

As $\mathcal{K}_{\mathcal{H}}$ is the inverse under composition of $\mathcal{G}_{\mathcal{H}}$, instead of minimizing over y, we will minimize over $x > c$ where $y = \mathcal{G}_{\mathcal{H}}(x)$, and we also observe that $\mathcal{G}_{\mathcal{H}}(x) = \mathcal{G}_H(x)$. Overall this leads us to

$$\alpha_2\left(\mathop{\mathbf{E}}_{\mathbf{P}} \chi_x\left(\mathbf{R}_{\mathbf{P}}\right)\right) \leq \min_{x > c} \mathcal{R}_H(x), \tag{43}$$

where we recall that $\alpha_k(p)$ is the k-th largest root of the polynomial $p(x)$.

For proving the existence of a graph achieving the same bound as the expected characteristic polynomial, we apply Lemma 4.9 with $k = 2$, as we see that $\mathbf{E}_{\mathbf{P}}\, \chi_x\left(\mathbf{R}_{\mathbf{P}}\right)$ aligns with the RHS of Equation (15). $\square$

6.2 Lower bound

Definition 6.2 ("Bad" graphs for the replacement product)**.** *We say that a vertex-transitive (d, c)-graph H is* bad *if the number of connected components in H is at least $\frac{d}{2}$. We say that a graph is good, being vertex transitive, if it is not bad.*

We note that bad graphs for the replacement product must be perfect matching graphs with c_1 parallel edges for every match, along with $c - c_1$ self-loops for each vertex, for some $c_1 \in \{0, \ldots, c\}$. The replacement product of any graph G with a bad graph is degenerate in some sense, and is amenable to a straightforward analysis. The replacement product with matching graphs (with self loops) result in a union of disjoint cycles (with self loops), and the replacement product with graphs with self loops only result in perfect matching graphs with self loops.

Theorem 6.3. *Let H be a good vertex-transitive (d, c)-graph with $d \geq 3$ and $c \geq 2$. Then,* *for every (n,d)-graph G,*

$$\lambda(G \textcircled{r} H) \geq \min_{x>c} \mathcal{R}_H(x) - o_n(1).$$

To achieve a lower bound to a expansion of the replacement product with H, we stick to our proof of the lower bound for the Zig-Zag product.

6.2.1 Symbolic analysis

Let H be a good vertex-transitive (d,c)-graph with $d \geq 3$ and $c \geq 2$. First, we look at the replacement product of the infinite d-ary tree and H, denoted as $\mathcal{T}_d \textcircled{r} H$. Our goal with this object is to lower bound the number of cycles of every length that originate at a fixed root vertex r to itself. Recall that the edges adjacent to each vertex consist of a single edge that is induced by the rotation map of G and c edges which are induced by H. To simplify the task of lower bounding the number of cycles, we truncate the product graph $\mathcal{T}_d \textcircled{r} H$. In this case, the truncation is rather simple: simply remove the edge that touches r that comes from the rotation map of G. Denote the truncated version by X.

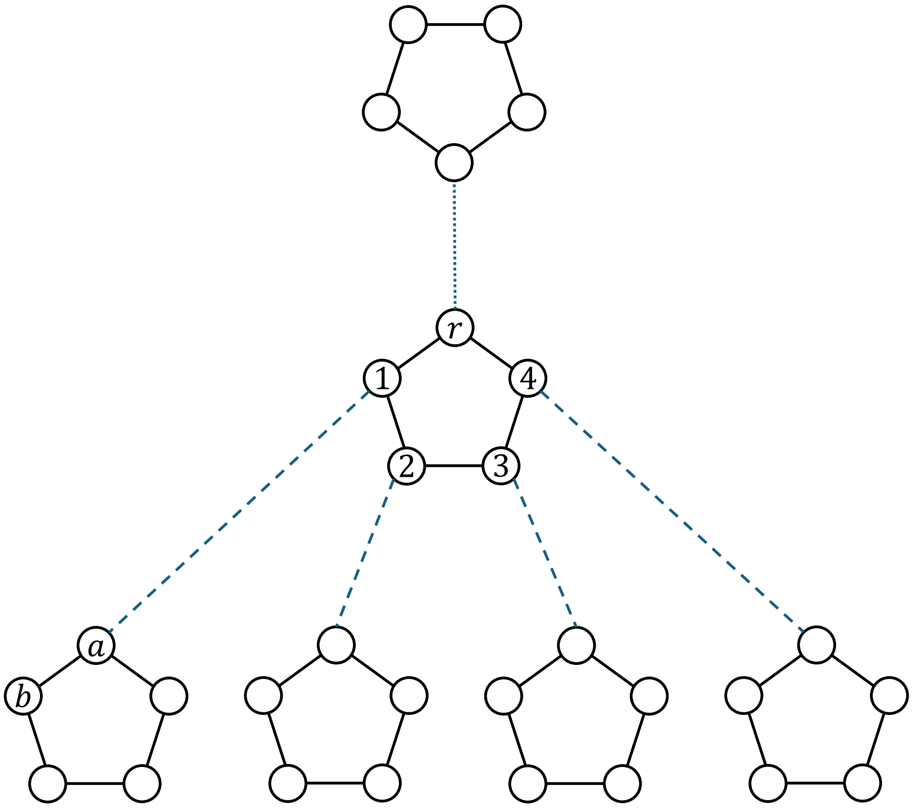

Figure 4: The replacement product graph $\mathcal{T}_5 \textcircled{r} C_5$. The black solid edges are those of C_5. The blue dashed edges and the dotted edge come from the rotation map of $\mathcal{T}_5$. In the truncated graph X, the dotted edge is truncated.

Fixing an arbitrary vertex v in H, and having r, the root of X fixed as before, we define three combinatorial classes:

1. $\mathcal{S}_H$: Cycles in H, originating at the vertex v, returning to v only upon completing the cycle (excluding the empty cycle), with corresponding generating function $S_H(z)$.
2. $\mathcal{D}_H$: Cycles in $\mathcal{S}_H$, where the last step is omitted, with the corresponding generating function $D_H(z)$.
3. $\mathcal{S}_X$: Cycles in X, originating at the vertex r, returning to r only upon completing the cycle (excluding the empty cycle), with the corresponding generating function $S_X(z)$.

Note that $\mathcal{D}_H$ represents the paths from v to one of its neighbors that do not traverse v in the middle. By the definition of $\mathcal{D}_H$, it is clear that its generating function satisfies $D_H(z) = \frac{S_H(z)}{z}$. In a similar fashion to Claim 5.3, the relation between S_H and the Cauchy transform is given by $S_H(z) = 1 - \frac{1}{\frac{1}{z}\mathcal{G}_H\left(\frac{1}{z}\right)}$.

Although we think of edges that come from H and edges that come from $\mathcal{T}_d$ as equal, it will be didactic to denote them differently. The atomic symbolic variable $\mathcal{Z}$ will capture an H-edge and the atomic symbolic variable $\mathcal{U}$ will capture a $\mathcal{T}_d$-edge.

Lemma 6.4. *The class $\mathcal{S}_X$ satisfies the symbolic relation*

$$\mathcal{S}_X = \mathcal{D}_H(\mathcal{Z} \times \mathsf{SEQ}(\mathcal{U} \times \mathsf{SEQ}(\mathcal{S}_X) \times \mathcal{U})) \times \mathcal{Z}.$$

Proof. In order to close a cycle from r to itself, one must close a cycle in the top-level copy of r (i.e., the copy of H which r is a part of). Each step in that copy of H (except for the last one) can then be replaced by performing the step and then adding a detour of a specific form that we will describe next. The description so far gives us the symbolic "template" for the formula

$$\mathcal{S}_X = \mathcal{D}_H(\mathcal{Z} \times \cdots) \times \mathcal{Z}, \tag{44}$$

where the latter $\mathcal{Z}$ describes the last step which cannot be replaced by a detour and $\mathcal{D}_H$ already has a $\mathcal{Z}$ in its input because we replace each step by a pair consisting of a step in H and a detour. For describing the detour, let us start with an example. Consider the cycle $(r, 1, 2, 3, 4, r)$ in Figure 4. The first step $r \to 1$ can be replaced, for example by $r \to 1 \to a \to 1$. In general, the alternate route can go down the blue dashed edge $(1, a)$, perform an arbitrary cycles from a to itself that does not use the edge $(1, a)$ and then return to vertex 1. Note that the cycles from a to itself that do not use the edge $(1, a)$ are exactly captured by the symbolic expression $\mathsf{SEQ}(\mathcal{S}_X)$. This specification of alternate route generalizes to the general case, and is captured by the symbolic expression $\mathcal{U} \times \mathsf{SEQ}(\mathcal{S}_X) \times \mathcal{U}$. We now observe that a detour can queue several such alternate routes, which yields the symbolic representation of the detour $\mathsf{SEQ}(\mathcal{U} \times \mathsf{SEQ}(\mathcal{S}_X) \times \mathcal{U})$. Plugging this back in Equation (44), we get the desired result. $\square$

From the symbolic relation we obtain a functional equation. As we count H-steps the same as $\mathcal{T}_d$-steps, we assign the trivial generating functions $Z(z) = z = U(z)$. Lemma 6.4 then yields the functional equation

$$S_X(z) = z \cdot D_H\left(z \cdot \frac{1}{1 - \frac{z^2}{1-S_X(z)}}\right) = z \cdot D_H\left(\frac{z(1-S_X(z))}{1-S_X(z)-z^2}\right). \tag{45}$$

$S_X(z)$ is now amenable to analysis using Theorem 4.17. In order to invoke this theorem, we first need to prove that $S_X(z)$ belongs to the smooth implicit-function schema. In particular, we need to solve the characteristic equation induced by $S_X(z)$, which is given by

$$w = z \cdot D_H\left(\frac{z(1-w)}{1-w-z^2}\right), \tag{46}$$

$$1 = z \cdot D'_H\left(\frac{z(1-w)}{1-w-z^2}\right) \cdot \frac{z^3}{(1-w-z^2)^2}. \tag{47}$$

6.2.2 Solving the characteristic system and minimizing $\mathcal{R}_H$

In this section we will show that $\mathcal{R}_H(x)$ must attain a minimum in the range $x > c$, and then relate its minimizer to the solution of the characteristic system

Lemma 6.5. *The function $\mathcal{R}_H(x)$ has a minimum in the domain $x > c$.*

Proof. It suffices to show that $\mathcal{R}'_H(x)$ is negative near $x = c$ and that $\mathcal{R}_H(x) \to \infty$ as $x \to \infty$.

$$\mathcal{R}'_H(x) = 1 + \frac{\mathcal{G}'_H(x)}{2\mathcal{G}^2_H(x)} \cdot \left(1 - \frac{1}{\sqrt{4\mathcal{G}^2_H(x)+1}}\right). \tag{48}$$

As x approaches c, $G^2_H(x)$ diverges to ∞, and so the term in the parentheses approaches 1. Similarly to Equation (26), the ratio $\frac{\mathcal{G}'_H(x)}{\mathcal{G}^2_H(x)}$ approaches $-\frac{d}{m_1}$ as x approaches c, where m_1 is the multiplicity of the eigenvalue c in H. For good graphs (in the sense of Definition 6.2), we get that $\mathcal{R}'_H(x)$ approaches $1 - \frac{d}{2m_1} < 0$.

As for the limit of $\mathcal{R}_H(x)$ as x goes to ∞, note that $\mathcal{R}_H(x) > x$ and hence diverges to ∞. These two facts together with the continuity of the derivative of $\mathcal{R}_H(x)$ in the domain (c, ∞) imply the existence of a minimum of $\mathcal{R}_H(x)$ in the same domain. □

To solve the characteristic system, it will be convenient to introduce a new variable

$$x(z, w) = \frac{z(1-w)}{1-w-z^2} = \frac{\frac{1}{z}}{\frac{1}{z^2} - \frac{1}{1-w}}. \tag{49}$$

Expressing w in terms of x and z, we get the reverse relation

$$w(z,x) = \frac{x - z - xz^2}{x - z} = 1 - \frac{z}{\frac{1}{z} - \frac{1}{x}}. \tag{50}$$

This gives us a characteristic system in the variables z and x instead of z and w. Working out the substitution, one gets arrived at the system

$$1 = \left(\frac{1}{z} - \frac{1}{x}\right) \cdot \left(\frac{1}{z} - D_H(x)\right), \tag{51}$$

$$1 = \left(\frac{1}{z} - \frac{1}{x}\right)^2 \cdot (x^2 D_H'(x)). \tag{52}$$

As we believe (and later on prove), the lower and upper bounds coincide, and together with Theorem 4.17 and Theorem 6.1, choosing z_0 such that $\frac{1}{z_0} = \min_{y>c} \mathcal{R}_H(y)$ is a good guess to start with. As it turns out, the minimizer $\arg\min_{y>c} \mathcal{R}_H(y)$ also plays a role in the solution of the characteristic equation.

Lemma 6.6. *Assume that $\mathcal{R}_H'(t) = 0$. Then the characteristic system in the variables z and x, Equations (51) and (52), has a solution*

$$(z_0, x_0) = \left(\frac{1}{\mathcal{R}_H(t)}, \frac{1}{t}\right).$$

Proof. For Equation (51) we simply have to recall that

$$D_H(x) = \frac{S_H(x)}{x} = \frac{1}{x} - \frac{1}{\mathcal{G}_H(\frac{1}{x})}, \tag{53}$$

and so, recalling Equation (40),

$$\begin{aligned}
\frac{1}{z_0} - D_H(x_0) &= \mathcal{R}_H(t) - D_H\left(\frac{1}{t}\right) \\
&= t + \frac{\sqrt{1 + 4\mathcal{G}_H(t)^2} - 1}{2\mathcal{G}_H(t)} - \left(t - \frac{1}{\mathcal{G}_H(t)}\right) \\
&= \frac{\sqrt{1 + 4\mathcal{G}_H(t)^2} + 1}{2\mathcal{G}_H(t)}.
\end{aligned}$$

Additionally, we have that

$$\frac{1}{z_0} - \frac{1}{x_0} = \mathcal{R}_H(t) - t = \frac{\sqrt{1 + 4\mathcal{G}_H(t)^2} - 1}{2\mathcal{G}_H(t)}. \tag{54}$$

Plugging these into Equation (51), we obtain

$$\left(\frac{1}{z_0} - \frac{1}{x_0}\right) \cdot \left(\frac{1}{z_0} - D_H(x_0)\right) = \frac{\sqrt{1+4\mathcal{G}_H(t)^2} - 1}{2\mathcal{G}_H(t)} \cdot \frac{\sqrt{1+4\mathcal{G}_H(t)^2} + 1}{2\mathcal{G}_H(t)} = 1,$$

which means that (z_0, x_0) is a valid solution to Equation (51). As for Equation (52), working out the derivatives, we obtain

$$\mathcal{R}'_H(t) = 0 \quad \Longrightarrow \quad \mathcal{G}'_H(t) = \frac{2\mathcal{G}_H^2(t)}{\frac{1}{\sqrt{1+4\mathcal{G}_H^2(t)}} - 1}.$$

Expressing D'_H in terms of $\mathcal{G}_H$, using Equation (53), we have the relation

$$D'_H(x) = -\frac{1}{x^2}\left(1 + \frac{\mathcal{G}'_H\left(\frac{1}{x}\right)}{\mathcal{G}_H^2\left(\frac{1}{x}\right)}\right),$$

and so

$$x_0^2 D'_H(x_0) = -\left(1 + \frac{2}{\frac{1}{\sqrt{1+4\mathcal{G}_H^2(t)}} - 1}\right) = \frac{\sqrt{1+4\mathcal{G}_H^2(t)} + 1}{\sqrt{1+4\mathcal{G}_H^2(t)} - 1}.$$

Furthermore, by Equation (54),

$$\left(\frac{1}{z_0} - \frac{1}{x_0}\right)^2 = \frac{\left(\sqrt{1+4\mathcal{G}_H(t)^2} - 1\right)^2}{4\mathcal{G}_H^2(t)}.$$

Plugging these together into Equation (52), noting that,

$$\frac{\left(\sqrt{1+4\mathcal{G}_H(t)^2} - 1\right)^2}{4\mathcal{G}_H^2(t)} \cdot \frac{1 + \sqrt{1+4\mathcal{G}_H^2(t)}}{1 - \sqrt{1+4\mathcal{G}_H^2(t)}} = 1,$$

we see that our chosen point (z_0, x_0) satisfies Equation (52), and hence is a solution to the system. □

Corollary 6.7. *Assume that $\mathcal{R}'_H(t) = 0$. The characteristic system in the variables z, w has the solution*

$$(z_0, w_0) = \left(\frac{1}{\mathcal{R}_H(t)}, 1 - \frac{1}{\mathcal{R}_H(t)(\mathcal{R}_H(t) - t)}\right),$$

which lies in the positive quadrant.

Proof. Let (z_0, x_0) be the solution to the characteristic system in the variables z and x obtained in Lemma 6.6. The fact that (z_0, w_0) solves the characteristic equation in the variables z and w is immediate from Lemma 6.6 and the relation $w = 1 - \frac{z}{\frac{1}{z} - \frac{1}{x}}$ given in

Equation (49). It is easy to verify that $\mathcal{R}_H(x) > x$ for all $x > c$. Hence,

$$x_0 = \frac{1}{t} < \frac{1}{\mathcal{R}_H(t)} < z_0.$$

By Equation (49),

$$w_0 = w(z_0, x_0) = \frac{x_0 - z_0 - x_0 z_0^2}{x_0 - z_0}.$$

As $0 < x_0 < z_0$, it is easy to verify that both the numerator and the denominator are negative, and hence w_0 is positive. □

Lemma 6.8. *The following conditions are equivalent:*

1. $t = \min_{y>c} \mathcal{R}_H(y)$,
2. *The point* $\left(\frac{1}{\mathcal{R}_H(t)}, 1 - \frac{1}{\mathcal{R}_H(t)(\mathcal{R}_H(t)-t)}\right)$ *is the solution to the characteristic system defined by* $S_X(z) = P(z, S_X(z))$ *inside the radius of convergence of* $P(z,w)$*, in the positive quadrant.*

Proof. As t is a local minimizer of $\mathcal{R}_H(y)$, $\mathcal{R}'_H(t)$ must vanish. We conclude, by Corollary 6.7, that the point $\left(\frac{1}{\mathcal{R}_H(t)}, 1 - \frac{1}{\mathcal{R}_H(t)(\mathcal{R}_H(t)-t)}\right)$ is a solution to the characteristic system within the radius of convergence of $P(z,w)$. For the other direction recall that the solution to the characteristic system in the positive quadrant inside the radius of convergence is guaranteed to be unique. Finally, note that the function $p(y) \triangleq \left(\frac{1}{\mathcal{R}_H(y)}, \frac{1}{\mathcal{R}_H(y)(\mathcal{R}_H(y)-y)}\right)$ is injective, proving the equivalence between the two conditions. □

Corollary 6.9. *There is a unique point* $t \in (c, \infty)$ *such that* $\mathcal{R}'_H(t) = 0$*. In* t*,* $\mathcal{R}_H(y)$ *attains its global minimum in that domain.*

Proof. By Theorem 4.17, the solution (z_0, w_0) to the characteristic system is unique. Lemma 6.8 then concludes the proof. □

6.2.3 Coefficients extraction

To extract the coefficients from $S_X(z)$, we want to invoke Theorem 4.17 with $y(z) = S_X(z)$ and $P(z,w) = z \cdot D_H\left(\frac{z(1-w)}{1-w-z^2}\right)$. We already reduced its characteristic system to the minimization of $\mathcal{R}_H(y)$. Denote $t = \arg\min_{y>c} \mathcal{R}_H(y)$. Let

$$(z_0, w_0) = \left(\frac{1}{\mathcal{R}_H(t)}, 1 - \frac{1}{\mathcal{R}_H(t)(\mathcal{R}_H(t) - t)}\right)$$

be the solution to the characteristic system.

Lemma 6.10. *There exists $\varepsilon > 0$ such that $P(z,w)$ is analytic in the domain*

$$\{(z,w) : |z| < z_0 + \varepsilon, |w| < w_0 + \varepsilon\}.$$

Proof. To prove the lemma we use the fact that the analyticity of functions is closed under composition. Recall that $P(z,w) = z \cdot D_H(x(z,w))$. We will first show that the radius of convergence of D_H is $\frac{1}{c}$. We will then show that there exists $\varepsilon > 0$ such that $x(z,w)$ is analytic in the domain $|z| < z_0 + \varepsilon,\ |w| < w_0 + \varepsilon$, and that in this domain $|x(z,w)| < \frac{1}{c}$.

As D_H has non-negative coefficients, per Equation (18), to compute its radius of convergence it suffices to find its first singularity along the non-negative part of the real axis. Recall that by Equation (53), $D_H(x) = \frac{1}{x} - \frac{1}{\mathcal{G}_H(\frac{1}{x})}$. Note that as the smallest element in the class $\mathcal{D}_H$ is of size 1, and so $D_H(0) = 0$. In particular it does not have a singularity at 0. The singularities of D_H must thus come from the singularities of $\frac{1}{\mathcal{G}_H(\frac{1}{x})}$. Similarly to Claim 5.3, the radius of convergence of D_H must thus come from the zeros of $\mathcal{G}_H(\frac{1}{x})$. Since $\mathcal{G}_H(y) > 0$ for all $y > c$, the radius of convergence of $D_H(x)$ is lower bounded by $\frac{1}{c}$. As $x_0 = \frac{1}{t} < \frac{1}{c}$, there exists $\varepsilon_1 > 0$ such that $D_H(x)$ is analytic for all x such that $|x| < x_0 + \varepsilon_1$.

To analyze the singularities of x as a function of z and w, we recall Equation (50),

$$|x(z,w)| = \frac{\left|\frac{1}{z}\right|}{\left|\frac{1}{z^2} - \frac{1}{1-w}\right|} \leq \frac{\left|\frac{1}{z}\right|}{\left|\frac{1}{z}\right|^2 - \left|\frac{1}{1-w}\right|} \leq \frac{\frac{1}{|z|}}{\frac{1}{|z|^2} - \frac{1}{1-|w|}}.$$

Thus, it suffices to prove that $x(z,w) < \frac{1}{c}$ for real non-negative inputs within our domain $|z| < z_0 + \varepsilon,\ |w| < w_0 + \varepsilon$. For this we first note that

$$\begin{aligned} 1 - w_0 - z_0^2 &= 1 - \left(1 - \frac{1}{\mathcal{R}_H(t)(\mathcal{R}_H(t) - t)}\right) - \frac{1}{\mathcal{R}_H(t)^2} \\ &= \frac{1}{\mathcal{R}_H(t)(\mathcal{R}_H(t) - t)} - \frac{1}{\mathcal{R}_H(t)^2} \\ &> 0, \end{aligned}$$

where the inequality follows from the fact $\mathcal{R}_H(t) > t$. Therefore, there exists $\varepsilon_2 > 0$ such that

$$1 - (w_0 + \varepsilon_2) - (z_0 + \varepsilon_2)^2 > 0.$$

We now claim that within the domain $0 < z < z_0 + \varepsilon_2,\ 0 < w < w_0 + \varepsilon_2$, the function

$x(z,w)$ is monotone increasing in z and in w. Indeed,

$$\frac{\partial x}{\partial z} = \frac{(1-w)(1-w+z^2)}{(1-w-z^2)^2} > 0,$$
$$\frac{\partial x}{\partial w} = \frac{z^3}{(1-w-z^2)^2} > 0.$$

As $x(z,w)$ is continuous within the aforementioned domain, it must be monotone increasing. Thus, it suffices to prove that $x(z_0,w_0) < \frac{1}{c}$, which indeed holds as $x(z_0,w_0) = x_0 = \frac{1}{t} < \frac{1}{c}$. So there exists $\varepsilon > 0$ such that within the domain $|z| < z_0+\varepsilon,\ |w| < w_0+\varepsilon$, we have that $|x(z,w)| < \frac{1}{c}$ as desired. Altogether, the singularities of $P(z,w)$ consist of the singularities of $x(z,w)$ together with the points in which $x(z,w)$ hits a singularity of D_H. These two cases are never obtained in the domain $|z| < z_0+\varepsilon,\ |w| < w_0+\varepsilon$ and so $P(z,w)$ is analytic in that domain. □

Lemma 6.11. *The coefficients of $P(z,w)$ satisfy condition 2 in Definition 4.16.*

Proof. We follow the symbolic derivation in order to prove the non-negativity of the coefficients. w, as a power series with a single non-zero coefficient, clearly has non-negative coefficients. $\frac{1}{1-w}$, as it represents sequences of the element represented by a series with non-negative coefficients, also has only non-negative coefficients. Multiplying everything by z^2 does not affect the non-negativity of the coefficients, but only shifts them. Taking another sequence of the obtained objects and using the same argument, we have that $\frac{1}{1-\frac{z^2}{1-w}}$ has non-negative coefficients. Again, multiplying everything by z dose not affect the non-negativity. So

$$\frac{z}{1-\frac{z^2}{1-w}} = \frac{z(1-w)}{1-w-z^2} = x(z,w)$$

has non-negative coefficients. Finally, as it counts paths, D_H has only non-negative coefficients. When we plug a non-negative series into it we must obtain a non-negative series. Hence, $D_H\left(\frac{z(1-w)}{1-w-z^2}\right)$ has non-negative coefficients. Finally, another multiplication by z shifts the coefficients and leaves them non-negative, and $P(z,w)$ must have only non-negative coefficients. Moving on to the next bullet, $p_{0,0} = P(0,0) = 0\cdot D_H(0) = 0$, as desired. For the third bullet, we derive P w.r.t w, which gives

$$P_w(z,w) = z\cdot D'_H(x(z,w))\frac{z^3}{(1-w-z^2)^2}. \tag{55}$$

Plugging $z = w = 0$ we obtain $p_{0,1} = 0 \neq 1$. For the last bullet we need to show that P is not affine in w. To this end, it suffices to show that for an arbitrary value of z, P is not affine in w. Plug $z = \frac{1}{2c}$ in Equation (55). We claim that this is a strictly monotone increasing function of w (and not constant as in derivatives of affine functions). Indeed,

$x = x(z,w)$ is a strictly monotone increasing function of w near the point $(z,w) = \left(\frac{1}{2c}, 0\right)$. D_H is strictly monotone increasing in its radius of convergence, which is where x lies when $z = \frac{1}{2c}$ and w near 0. The outer term $\frac{z^4}{(1-w-z^2)^2}$ is also strictly monotone increasing in w near 0. So the derivative of P near 0 is not constant which means that P is not affine in w, as desired. $\square$

We are now ready to prove Theorem 6.3. First, we have that for every (n,d)-graph G,

$$c_\ell(G \textcircled{z} H)^{\frac{1}{\ell}} \geq c_\ell(\mathcal{T}_d \textcircled{z} H)^{\frac{1}{\ell}} \geq c_\ell(X)^{\frac{1}{\ell}} \geq \left([z^\ell]S_X(z)\right)^{\frac{1}{\ell}}.$$

Second, Lemma 6.10, Lemma 6.11 and Corollary 6.7 together with Corollary 6.9, show that $S_X(z)$ belongs to the smooth implicit-function schema. Invoking Theorem 4.17, we get that

$$S_X \bowtie \mathcal{R}_H(t) = \min_{x>c} \mathcal{R}_H(x).$$

To conclude the proof of Theorem 6.3, the application of Lemma 4.5 then gives us

$$\lambda(G \textcircled{r} H) \geq (1 - o_n(1))\mathcal{R}_H(t) = \min_{x>c} \mathcal{R}_H(x) - o_n(1).$$

6.3 Universal lower bound depending on degree only

In a similar fashion to Section 5.2, we prove a lower bound on the expansion of $G \textcircled{r} H$ that only depends on its degree, and captures the best possible behavior we can expect from a replacement product.

Theorem 6.12. *Let H be a good vertex-transitive (d,c)-graph with $d \geq 3$ and $c \geq 3$. Then, for every (n,d)-graph G,*

$$\lambda(G \textcircled{r} H) \geq c + \frac{c}{c-1} - o_n(1). \tag{56}$$

For $c = 2$,

$$\lambda(G \textcircled{r} H) \geq 2\sqrt{2} - o_n(1) \approx 2.8284.$$

For $c \geq 3$, as the degree of $G \textcircled{r} H$ is $c+1$, we have that

$$\omega(G \textcircled{r} H) \geq \frac{c}{c+1},$$

meaning that the replacement product cannot produce good expanders, and the rate at which $\omega(G \textcircled{r} H)$ approaches 1 is inverse linear in c.

Proof. We follow the same lines of the proof of Theorem 5.12, where in this case we replace $\mathcal{G}_H$ in Theorem 6.3 with the Cauchy transform of the Kesten-McKay distribution, which is known from the free probability literature to be

$$\mathcal{G}_{\mathsf{km}}(x) = \frac{c\sqrt{x^2 - 4(c-1)} - x(c-2)}{2(x^2 - c^2)}.$$

Let

$$\varphi(x, y) = x + \frac{\sqrt{1 + 4y^2} - 1}{2y}.$$

We note that $\varphi(x, y)$ is monotone increasing in y and that $\mathcal{R}_H(x) = \varphi(x, \mathcal{G}_H(x))$, and we denote $\mathcal{R}_{\mathcal{T}_c}(x) = \varphi(x, \mathcal{G}_{\mathsf{km}}(x))$. By Equation (30) we get that $\mathcal{R}_H(x) \geq \mathcal{R}_{\mathcal{T}_c}(x)$ for all $x > c$.

For the $c = 2$ case we have that

$$\mathcal{R}_{\mathcal{T}_2}(x) = x + \frac{\sqrt{x^2 - 4}}{2} \cdot \left(\sqrt{\frac{x^2}{x^2 - 4}} - 1 \right),$$

having a global minimum at $x_0 = \frac{3}{\sqrt{2}}$ with value $\mathcal{R}_{\mathcal{T}_2}(x_0) = 2\sqrt{2}$.

For $c \geq 3$, unfortunately, we cannot deduce a nice formula similar to Equation (32) for $\mathcal{R}_{\mathcal{T}_c}(x)$. However it can be verified that $\mathcal{R}_{\mathcal{T}_c}(x)$ is increasing in (c, ∞), hence the minimal value is achieved at $\mathcal{R}_{\mathcal{T}_c}(c)$, which will be our lower bound. Note that c is a removable singularity for $\mathcal{G}_{\mathsf{km}}(x)$, so we are formally calculating $\lim_{x \to c^+} (\mathcal{R}_{\mathcal{T}_c}(x))$.

Denote $g_c = \mathcal{G}_{\mathsf{km}}(c)$. By a straightforward calculation, $g_c = \frac{c-1}{c(c-2)}$ for every $c \geq 3$, and we get that

$$\mathcal{R}_{\mathcal{T}_c}(c) = \varphi(c, g_c) = c + \frac{\sqrt{1 + 4g_c^2} - 1}{2g_c} = c + \frac{1}{c-1}.$$

$\square$

References

[AC88] N. Alon and F. R. K. Chung. Explicit construction of linear sized tolerant networks. In *Proceedings of the First Japan Conference on Graph Theory and Applications (Hakone, 1986)*, number 1-3, pages 15–19, 1988.

[BATS11] Avraham Ben-Aroya and Amnon Ta-Shma. A combinatorial construction of almost-Ramanujan graphs using the zig-zag product. *SIAM J. Comput.*, 40(2):267–290, 2011.

[CCMP23] Gil Cohen, Itay Cohen, Gal Maor, and Yuval Peled. Derandomized squaring: An analytical insight into its true behavior. In *Electronic Colloquium on Computational Complexity (ECCC)*, volume 183, 2023.

[CM23] Gil Cohen and Gal Maor. Random walks on rotating expanders. In *Proceedings of the 55th Annual ACM Symposium on Theory of Computing*, pages 971–984, 2023.

[Coh16] Michael B. Cohen. Ramanujan graphs in polynomial time. In *57th Annual IEEE Symposium on Foundations of Computer Science—FOCS 2016*, pages 276–281. IEEE Computer Soc., Los Alamitos, CA, 2016.

[Din07] Irit Dinur. The PCP theorem by gap amplification. *J. ACM*, 54(3):Art. 12, 44, 2007.

[FS09] Philippe Flajolet and Robert Sedgewick. *Analytic combinatorics*. Cambridge University Press, 2009.

[Gol05] Oded Goldreich. Bravely, moderately: A common theme in four recent results. In *Electronic Colloquium on Computational Complexity (ECCC)*, volume 98, page 4, 2005.

[HLW06] Shlomo Hoory, Nathan Linial, and Avi Wigderson. Expander graphs and their applications. *Bulletin of the American Mathematical Society*, 43(4):439–561, 2006.

[KMRZS17] Swastik Kopparty, Or Meir, Noga Ron-Zewi, and Shubhangi Saraf. High-rate locally correctable and locally testable codes with sub-polynomial query complexity. *J. ACM*, 64(2):Art. 11, 42, 2017.

[McK81] Brendan D McKay. The expected eigenvalue distribution of a large regular graph. *Linear Algebra and its Applications*, 40:203–216, 1981.

[MSS15] Adam W. Marcus, Daniel A. Spielman, and Nikhil Srivastava. Interlacing families II: Mixed characteristic polynomials and the Kadison-Singer problem. *Ann. of Math. (2)*, 182(1):327–350, 2015.

[MSS18] Adam W Marcus, Daniel A Spielman, and Nikhil Srivastava. Interlacing families IV: Bipartite ramanujan graphs of all sizes. *SIAM Journal on Computing*, 47(6):2488–2509, 2018.

[MSS22] Adam W. Marcus, Daniel A. Spielman, and Nikhil Srivastava. Finite free convolutions of polynomials. *Probab. Theory Related Fields*, 182(3-4):807–848, 2022.

[Nil91] A. Nilli. On the second eigenvalue of a graph. *Discrete Math.*, 91(2):207–210, 1991.

[NS06] Alexandru Nica and Roland Speicher. *Lectures on the combinatorics of free probability*, volume 335 of *London Mathematical Society Lecture Note Series*. Cambridge University Press, Cambridge, 2006.

[Rei08] Omer Reingold. Undirected connectivity in log-space. *J. ACM*, 55(4):Art. 17, 24, 2008.

[RV05] Eyal Rozenman and Salil Vadhan. Derandomized squaring of graphs. In *Approximation, randomization and combinatorial optimization*, volume 3624 of *Lecture Notes in Comput. Sci.*, pages 436–447. Springer, Berlin, 2005.

[RVW00] Omer Reingold, Salil Vadhan, and Avi Wigderson. Entropy waves, the zig-zag graph product, and new constant-degree expanders and extractors. In *Proceedings 41st Annual Symposium on Foundations of Computer Science*, pages 3–13. IEEE, 2000.

[TS17] Amnon Ta-Shma. Explicit, almost optimal, epsilon-balanced codes. In *Proceedings of the 49th Annual ACM SIGACT Symposium on Theory of Computing*, pages 238–251, 2017.

A The symbolic method: a brief overview

In this section we give a brief overview for the symbolic method. We believe that the unfamiliar reader will benefit from reading it. The symbolic method provides a technique to convert a specification of a combinatorial class by means of certain combinatorial constructs into a functional equation that is satisfied by its associated generating function. In more technical terms, a *combinatorial class* $\mathcal{A}$ consists of a collection of combinatorial objects paired with a designated *size* function $|\cdot| : \mathcal{A} \to \mathbb{N}$. The associated *(ordinary) generating function* for this class is the formal power series $A(z) = \sum_{a \in \mathcal{A}} z^{|a|} = \sum_{k \in \mathbb{N}} A_k z^k$, where A_k is the number of objects in $\mathcal{A}$ of size k, which we always assume is finite.

Set theoretic operators on the combinatorial classes reflect in their associated generating functions. For instance, when two combinatorial classes, $\mathcal{A}$ and $\mathcal{B}$, are combined in a disjoint union, denoted as $\mathcal{A} + \mathcal{B}$, the corresponding generating function becomes the sum of their individual generating functions, $A(z) + B(z)$. When considering the Cartesian product $\mathcal{A} \times \mathcal{B}$, it corresponds to the multiplication of their generating functions, $A(z) \cdot B(z)$. In this context, the size of an element (a, b) from $\mathcal{A} \times \mathcal{B}$ is given by $|a| + |b|$. This concept of the Cartesian product can be extended to multiple classes. Another valuable concept is the *sequence* of a class, denoted as $\mathsf{SEQ}(\mathcal{A})$. This represents the disjoint union of the Cartesian products across all finite lengths $n \geq 0$. The generating function for $\mathsf{SEQ}(\mathcal{A})$ is given by $\frac{1}{1-A(z)}$.

We make use of standard shorthand notations: For an integer $\ell \geq 1$ and a class $\mathcal{A}$, we let $\ell\mathcal{A}$ denote the sum of ℓ copies of $\mathcal{A}$. We similarly write $\mathcal{A}^\ell$ for the Cartesian product of ℓ copies of $\mathcal{A}$. The class denoted $\mathcal{Z}$ refers to the class containing a single element of size 1. Its generating function is, of course, z. The elements of size 1 in a combinatorial class $\mathcal{A}$ are called *atoms*, all of which are considered distinct. An element of size 0, denoted as ε, is called a *neutral object*.

For instance, the class of binary strings can be constructed as $\mathsf{SEQ}(\{0\} + \{1\})$ where both elements $0, 1$ in their corresponding sets are atoms. Note that we can also write the class more succinctly as $\mathsf{SEQ}(\mathcal{Z} + \mathcal{Z})$ or $\mathsf{SEQ}(2\mathcal{Z})$ as indeed $2\mathcal{Z}$ is a combinatorial class that consists of two atoms. For the purpose of counting elements, these descriptions are equivalent, or isomorphic, though the second is less informative. From this, the corresponding generating function is immediately obtained, $\frac{1}{1-2z} = \sum_{k=0}^{\infty} 2^k z^k$.

To give another example, consider the class of rooted trees where the sequence order of a node's children matters, meaning they are arranged from left to right. This class can be formulated using the recurrence $\mathcal{A} = \bullet \times \mathsf{SEQ}(\mathcal{A})$, where $\bullet$ symbolizes an atom denoting a node. In this context, the size function corresponds to the number of vertices in the tree. To elaborate, a tree consists of a node, contributing a size of 1, followed by a sequence of trees. The related generating function satisfies the functional equation $A(z) = \frac{z}{1-A(z)}$, or equivalently $A(z)^2 - A(z) + z = 0$. Using basic methods, it can be shown that the coefficients of $A(z)$ are the Catalan numbers.

B An elementary combinatorial proof of a weaker version of Theorem 5.12

We provide here a proof of a weaker version of the universal lower bound, namely the one expressed in Equation (33), using the trace method, without the usage of analytic combinatorics but rather based only on elementary combinatorial techniques. One may think of the main part of the proof as a count of cycles in $\mathcal{T}_d \textcircled{z} \mathcal{T}_c$, that is, the Zig-Zag operation with the graph H being the infinite c-regular tree.

Theorem B.1 gives a bound that is stronger than the Alon-Boppana bound, however the asymptotic dependency on the degree is similar.

Theorem B.1. *For every (n, d)-graph G and every (d, c)-graph H,*

$$\lambda(G \textcircled{z} H) \geq \frac{3\sqrt{3}}{2}\sqrt{c(c-1)} - o_n(1) \approx 2.5981\sqrt{c(c-1)}. \tag{57}$$

To establish the proof, we once again count the cycles in $\mathcal{T}_d \textcircled{z} H$ for applying the trace method. In a similar manner to the way cycles are counted in $\mathcal{T}_c$ using the Catalan numbers, we will reduce the count of cycles in $\mathcal{T}_d \textcircled{z} H$ to the *Fuss-Catalan numbers*. Note that the

Fuss-Catalan numbers played an important role in the proof of the lower bound on the spectral expansion of rotating expanders in [CCMP23], and our proof of Theorem B.1 is similar.

Definition B.2 (Fuss-Catalan numbers)**.** *For integers $p \geq 2$ and $k \geq 0$, the (p, k)* Fuss-Catalan number*, denoted as $C_k^{(p)}$, is given by*

$$C_k^{(p)} = \frac{1}{pk+1}\binom{pk+1}{k}.$$

Note that the ordinary Catalan numbers are obtained by setting $p = 2$. As for Catalan numbers, for every $p \geq 2$, Fuss-Catalan numbers also satisfy a recurrence relation which is given by

$$C_{k+1}^{(p)} = \sum_{a_1+a_2+\cdots+a_p=k} \prod_{i=1}^{p} C_{a_i}^{(p)}. \tag{58}$$

It is a known fact that

$$C_k^{(p)} = \left(\frac{1}{k}\right)^{\Theta(1)} \left(\frac{p^p}{(p-1)^{p-1}}\right)^k. \tag{59}$$

Proof of Theorem B.1. Consider a length-$2k$ cycle in $\mathcal{T}_d ⓩ H$, starting at an arbitrary vertex (v, i) which we consider as the root. By the breakdown of the steps described in Definition 4.4, it consists of $4k$ steps in H, where each step can be thought of as either a *forward* step or a *back* step. By forward and back we trace the distance from (v, i), the odd steps defining the distance of the *cloud* we are in: a forward odd step results in moving a cloud away, while a back odd step results in returning a cloud closer.

In order to enumerate this type of closed walks, we define an algorithm that uniquely defines such a walk:

1. Pick an integer $m \leq k$ such that a cycle is complete for the first time after $4m$ H-steps.

2. Pick an integer $\ell \leq m$ such that the walk got back to the original cloud for the first time at step $4\ell - 1$.

3. Recursively continue for the 3 gaps, which are $g_1 = 4\ell - 4$, $g_2 = 4m - 4\ell$, and $g_3 = 4k - 4m$.

Notice that $g_1 + g_2 + g_3 = 4(k-1)$, and moreover, for every choice of two forward steps, one of them has c options and the other $c - 1$. Hence the number of choices can be expressed recursively as

$$B_k = c(c-1) \cdot \sum_{a+b+c=k-1} B_a B_b B_c, \tag{60}$$

with the initial value $B_0 = 1$. This matches the recursive formula of Equation (58) for $C_k^{(3)}$.

Combining this with Equation (59), we have that $C_{2k}(G \circledcirc H) \bowtie (c(c-1))^k \cdot \left(\frac{3^3}{2^2}\right)^k$, completing the proof. □

C Proof of Lemma 4.11

Adopting similar notations to those in [MSS22], we denote by $\mathbb{P}(d)$ the set of real-rooted polynomials of degree exactly d with a positive leading coefficient. We denote by $\mathbb{P}^+(d)$ the subset that consists only of polynomials with non-negative roots. For proving Lemma 4.11, we need to extend Definition 4.7 by defining the convolution operation on polynomials of different degrees.

Definition C.1 (Multiplicative convolution; Definition 1.4 of [MSS22])**.** *Let*

$$a(x) = \sum_{i=0}^{d} x^{d-i}(-1)^i a_i,$$
$$b(x) = \sum_{i=0}^{d} x^{d-i}(-1)^i b_i$$

be polynomials of degree at most *d. The d-th multiplicative convolution, denoted as $a \boxtimes_d b$, is defined as the polynomial*

$$(a \boxtimes_d b)(x) = \sum_{i=0}^{d} x^{d-i}(-1)^i \frac{a_i b_i}{\binom{d}{i}}.$$

It is proven in [MSS22] that when $a(x)$ and $b(x)$ are real-rooted and of degree *exactly* d, Definition C.1 and Definition 4.7 coincide (note that for any polynomial $a(x) \in \mathbb{P}(d)$, one can define the matrix A to be a diagonal matrix with the roots of $a(x)$ appearing on the diagonal). When d is clear from context, we omit it from the notation. In this section, we prove the following result.

Lemma C.2 (Lemma 4.11 restated)**.** *Let $q(x) \in \mathbb{P}^+(d)$ and $p(x) \in \mathbb{P}(d)$ such that $p(x)$ has at least one positive root. Then, for every $y > 0$,*

$$\mathcal{N}_{p \boxtimes q}(y) \leq \frac{y}{y+1} \cdot \mathcal{N}_p(y)\mathcal{N}_q(y). \tag{61}$$

We denote by D the derivative operators of polynomials. Namely, the derivative of a polynomial $f(x)$ is denoted by $Df(x)$. For a real number $\alpha > 0$ we let U_α be the operator on polynomial space mapping f to $f - \alpha D f$. Note that Equation (61) is equivalent to

$$\mathcal{S}_{p \boxtimes q}(y) \leq \mathcal{S}_p(y) \cdot \mathcal{S}_q(y), \tag{62}$$

where $\mathcal{S}_p(y) = \frac{y}{y+1}\mathcal{N}_p(y)$.

The following claim follows by the exact same arguments as in [MSS22], but is stated for more general polynomials. For completeness, we reproduce the proof.

Claim C.3 (Base case; Lemma 4.8 in [MSS22])**.** *Let $\lambda \neq 0$. If $p(x) = (x-\lambda)^d$ and $q(x) \in \mathbb{P}(d)$, then for all $y > 0$,*

$$\mathcal{S}_{p \boxtimes q}(y) = \mathcal{S}_p(y) \cdot \mathcal{S}_q(y).$$

Proof. We first observe that since $p(x) = (x-\lambda)^d$, we can directly calculate that

$$\mathcal{S}_p(y) = \lambda. \tag{63}$$

By Definition 4.7, invoked with $\mathbf{A} = \lambda \mathbf{I}$, we see that $p \boxtimes q$ has the same roots as q, multiplied by λ. Particularly,

$$(p \boxtimes q)(x) = \lambda^d q\left(\frac{x}{\lambda}\right).$$

Let $q_\lambda(x)$ be the polynomial satisfying $q_\lambda(x) = q\left(\frac{x}{\lambda}\right)$. Notice that the $\mathcal{M}$-transform (recall Equation (17)) depends only on the roots of the corresponding polynomial, and hence is invariant under multiplication of the polynomial by a constant. Thus,

$$\mathcal{M}_{p \boxtimes q}(x) = \mathcal{M}_{q_\lambda}(x). \tag{64}$$

By a direct calculation, $\mathcal{M}_{q_\lambda}(x) = \mathcal{M}_q(\frac{x}{\lambda})$, and so applying $\mathcal{N}_q$ to both sides gives us

$$\mathcal{N}_q\left(\mathcal{M}_{q_\lambda}(x)\right) = \frac{x}{\lambda}.$$

This, together with Equation (64), implies that $\mathcal{N}_{p \boxtimes q}(y) = \lambda \cdot \mathcal{N}_q(y)$ is the inverse function of $\mathcal{M}_{p \boxtimes q}(x)$. Overall we get that

$$\mathcal{S}_{p \boxtimes q}(y) = \frac{y}{y+1}\mathcal{N}_{p \boxtimes q}(y) = \lambda \cdot \frac{y}{y+1}\mathcal{N}_q(y) = \lambda \cdot \mathcal{S}_q(y) = \mathcal{S}_p(y) \cdot \mathcal{S}_q(y),$$

where the last equality follows by Equation (63). $\square$

Claim C.4. *Let $p(x) \in \mathbb{P}(d)$. For a real number $a \geq d$, define*

$$u(x) = p(x) - \frac{x}{a}p'(x).$$

Then, $u(x)$ is real-rooted. Furthermore, if $a > d$ then

$$\mathsf{maxroot}\left(u(x)\right) > \mathsf{maxroot}\left(p(x)\right) \geq \mathsf{maxroot}\left(p'(x)\right).$$

Proof. It is easy to see that every root of $p(x)$ which equals 0 or has multiplicity at least 2 is a root of $u(x)$. Therefore, it suffices to prove the claim for $p(x)$ having d distinct roots $\alpha_1, \ldots, \alpha_d$ which differ from 0.

Note that $u(x)$ is of degree d in the case the $a > d$, and $d-1$ if $a = d$. Observe also that if $u(x) = 0$ then $p(x) \neq 0$, as $p(x)$ and $p'(x)$ cannot vanish at the same point per our assumption. Define

$$f(x) = \frac{u(x)}{p(x)} = 1 - \frac{1}{a}\sum_{i=1}^{d} \frac{x}{x - \alpha_i},$$

and note that the following hold:

1. $u(x) = 0 \iff f(x) = 0$.
2. $f(x)$ has a pole at every root of $p(x)$.
3. $\lim_{x\to\infty} f(x) = 1 - \frac{d}{a}$.
4. $\lim_{x\to-\infty} f(x) = 1 - \frac{d}{a}$, and $f(x)$ approaches the limit from above.

Let $\beta_1 < \cdots < \beta_r$ be the negative roots and $\gamma_1 < \cdots < \gamma_s$ be the positive roots of $p(x)$. By the observations above, $u(x)$ has a root smaller than β_1 (by 4, including the case $a = d$) and in between every two consecutive negative ones, hence at least r negative roots. It also has a root in between every two consecutive positive roots, and a root larger than γ_s if $a > d$ (proving the furthermore part of the statement), hence at least $s-1$ positive roots if $a = d$ and at least s positive roots if $a > d$, completing the proof. $\square$

We make use of the following lemmata, proven in [MSS22], which we state here without a proof.

Lemma C.5 (*Pinching*; Lemma 4.1 in [MSS22]). *For $d \geq 2$, let $p(x) \in \mathbb{P}(d)$ be monic, and write $p(x) = \prod_{i=1}^{d}(x - \lambda_i)$, where $\lambda_1 \geq \lambda_2 \geq \cdots \geq \lambda_d$. Assume further that $\lambda_1 > \lambda_k$ for some $k \in [d]$. Then, for every $\alpha > 0$ there exist $\mu, \rho \in \mathbb{R}$ and two polynomials $\widetilde{p}(x) \in \mathbb{P}(d)$ and $\widehat{p}(x) \in \mathbb{P}(d-1)$ such that*

1. *$p(x) = \widehat{p}(x) + \widetilde{p}(x)$,*
2. *$\widetilde{p}(x) = (x - \mu)^2 \prod_{i\neq 1,k}(x - \lambda_i)$,*
3. *$\widehat{p}(x) = (x - \rho) \prod_{i\neq 1,k}(x - \lambda_i)$,*
4. *$\mathsf{maxroot}(U_\alpha p) = \mathsf{maxroot}(U_\alpha \widetilde{p}) = \mathsf{maxroot}(U_\alpha \widehat{p})$, and*

5. $\rho > \lambda_1 > \mu > \lambda_k$. In particular, if $d \geq 3$ then $\widehat{p}$ has at least two distinct roots.

Lemma C.6 (Lemma 4.2 in [MSS22]). *Let $f(x), g(x), h(x)$ be real-rooted polynomials with positive leading coefficients such that $f = g + h$. Then,*

$$\mathsf{maxroot}(f) \leq \max\left\{\mathsf{maxroot}(g), \mathsf{maxroot}(h)\right\},$$

with equality if and only if

$$\mathsf{maxroot}(f) = \mathsf{maxroot}(g) = \mathsf{maxroot}(h).$$

Lemma C.7 (Degree reduction; Lemma 4.9 in [MSS22]). *For $p(x) \in \mathbb{P}(d)$ and $q(x) \in \mathbb{P}(k)$, where $k < d$, it holds that*

$$p(x) \boxtimes_d q(x) = \left(p(x) - \frac{x}{d}Dp(x)\right) \boxtimes_{d-1} q(x).$$

(Note that $p(x) - \frac{x}{d}Dp(x)$ is a polynomial of degree $d-1$.)

Lemma C.8 (Lemma 4.10 in [MSS22]). *For $q(x) = (x-\lambda)^d$,*

$$(xD - d)q(x) = \lambda d(x-\lambda)^{d-1}.$$

For $q(x) \in \mathbb{P}^+(d)$, it holds that $(xD-d)q(x) \in \mathbb{P}^+(d-1)$, and

$$\mathsf{maxroot}(q(x)) \geq \mathsf{maxroot}\left((xD-d)q(x)\right).$$

Note that $(xD-d)q(x)$ is indeed real-rooted by Claim C.4. In the special case $q(x) \in \mathbb{P}^+(2)$ with distinct roots, strict inequality holds, namely,

$$\mathsf{maxroot}(q(x)) > \mathsf{maxroot}\left(xq'(x) - 2q(x)\right).$$

Moving forward, it will prove useful to express the $\mathcal{N}$-transform at a given point w as the maximal root of a polynomial. By definition, $\mathcal{M}_p(x) = x\mathcal{G}_p(x) - 1 = \frac{x}{d} \cdot \frac{p'(x)}{p(x)} - 1$, therefore we can write

$$\mathcal{M}_p(x) = w \quad \iff \quad \left(1 - \frac{xD}{d(w+1)}\right)p(x) = 0.$$

For $w > 0$, we define the operator V_w by

$$V_w p(x) = \left(1 - \frac{xD}{d(w+1)}\right)p(x),$$

and note that $V_w p(x)$ is real-rooted by Claim C.4. As $\mathcal{N}_p(w)$ is defined to be the max-inverse of $\mathcal{M}_p(x)$, we can write

$$\mathcal{N}_p(w) = \mathsf{maxroot}\left(V_w p(x)\right). \tag{65}$$

Claim C.9. *Let $d \geq 2$ and let $p(x) \in \mathbb{P}(d)$ be a polynomial with a positive leading coefficient and at least one positive root. Then, for every $w > 0$,*

$$\mathsf{maxroot}\left(V_w p(x)\right) > \mathsf{maxroot}\left(p(x)\right).$$

Proof. Let $r(x) = \frac{xp'(x)}{d(w+1)}$, so that

$$v(x) \triangleq V_w p(x) = p(x) - r(x).$$

Denote $\alpha = \mathsf{maxroot}\left(p(x)\right)$. We notice that if p_0 is the leading coefficient in $p(x)$, then $r_0 = \frac{p_0}{w+1}$ is the leading coefficient of $r(x)$. In particular, $r_0 < p_0$, and thus $v(x) > 0$ as $x \to \infty$. Therefore, it suffices to prove that there exists $\beta \geq \alpha$ such that $v(\beta) < 0$.

It is clear that the roots of $r(x)$ are those of $p'(x)$ and 0. It is also known that the roots of $p(x)$ and its derivative interlace, and so $\mathsf{maxroot}\left(r(x)\right) \leq \alpha$, and as $r_0 > 0$, we know that $r(x) > 0$ for $x > \alpha$. In the case that α is a simple root of $p(x)$, by interlacing we have that $r(\alpha) > 0$ and therefore $v(\alpha) < 0$. Otherwise, assume that $m \geq 2$ is the multiplicity of α as a root of $p(x)$. Then, as $\varepsilon \to 0^+$, $p(\alpha+\varepsilon) \approx \varepsilon^m$ whereas $r(\alpha+\varepsilon) \approx \varepsilon^{m-1}$, and therefore for $\varepsilon > 0$ small enough, we get $v(\alpha+\varepsilon) < 0$. ☐

Corollary C.10 (Similar to Corollary 4.11 in [MSS22], but for more general real-rooted polynomials)**.** *Let $w > 0, d \geq 2$, and let $p(x) \in \mathbb{P}(d)$ have at least two distinct roots, at least one of which positive. Then, there exist $\widetilde{p}(x) \in \mathbb{P}(d)$ and $\widehat{p}(x) \in \mathbb{P}(d-1)$ such that*

1. $p(x) = \widehat{p}(x) + \widetilde{p}(x)$,
2. $\mathsf{maxroot}(\widetilde{p}) \leq \mathsf{maxroot}(p)$,
3. $\mathsf{maxroot}(V_w p) = \mathsf{maxroot}(V_w \widetilde{p}) = \mathsf{maxroot}(V_w \widehat{p})$, *and*
4. *if $d \geq 3$, $\widehat{p}(x)$ has at least two distinct roots.*

Proof. Let $t = \mathsf{maxroot}(V_w p)$, and set $\alpha = \frac{t}{d(w+1)}$. By Claim C.9, we know that $t > 0$ and hence $\alpha > 0$, and so $U_\alpha p(x)$ is well-defined. We have that

$$U_\alpha p(x) - V_w p(x) = \frac{x-t}{d(w+1)} p'(x), \tag{66}$$

and therefore t is a root of $U_\alpha p(x)$. Furthermore, by Claim C.4, $t > \mathsf{maxroot}(p')$ and hence for every $x > t$ it holds that $p'(x) > 0$, thus $U_\alpha p(x) > V_w p(x)$. As a consequence, $\mathsf{maxroot}(U_\alpha p) = t$.

Let $\widetilde{p}(x) \in \mathbb{P}(d)$ and $\widehat{p}(x) \in \mathbb{P}(d-1)$ be the polynomials guaranteed to exist by Lemma C.5. Then,

$$\mathsf{maxroot}(U_\alpha \widetilde{p}) = \mathsf{maxroot}(U_\alpha \widehat{p}) = t. \tag{67}$$

We wish to also show that $t = \mathsf{maxroot}(V_w \widehat{p})$, and a similar argument will hold for $V_w \widetilde{p}$. As before, note that

$$U_\alpha \widehat{p}(x) - V_w \widehat{p}(x) = \frac{x-t}{d(w+1)} \widehat{p}'(x). \tag{68}$$

Assume $s > t$ is the maximal root of $V_w \widehat{p}$. Note that by Equation (67) together with Claim C.4, $s > \mathsf{maxroot}(\widehat{p}')$, and as $\widehat{p}'(x)$ and $U_\alpha \widehat{p}(x)$ have positive leading coefficients, $\widehat{p}'(s) > 0$ and $U_\alpha \widehat{p}(s) > 0$. Hence substituting $x = s$, we get a positive number on the LHS of Equation (68) and a negative one on its RHS, in contradiction. $\square$

Proof of Lemma C.2. We prove a stronger version of the theorem by induction on the degree d. Equation (62) holds for $d = 1$ by Claim C.3, which also covers the case that one of $p(x)$ or $q(x)$ having exactly one root. For the case that both $p(x)$ and $q(x)$ have at least two distict roots we prove that a strict inequality holds, namely,

$$\mathsf{maxroot}\left(V_w(p \boxtimes_d q)\right) < \frac{w}{w+1} \mathsf{maxroot}(V_w p) \cdot \mathsf{maxroot}(V_w q). \tag{69}$$

Fix $q(x) \in \mathbb{P}^+(d)$ with at least two distinct roots, and for $p(x) \in \mathbb{P}(d)$ define

$$\phi(p) \triangleq \frac{w}{w+1} \mathsf{maxroot}(V_w p) \cdot \mathsf{maxroot}(V_w q) - \mathsf{maxroot}(V_w(p \boxtimes_d q)).$$

Assume by contradiction that there exists $p(x) \in \mathbb{P}(d)$ such that $\phi(p) \leq 0$. Let $[-R, R]$ be an interval containing all roots of $p(x)$. Since $[-R, R]^d$ is a compact set and $\phi(p)$ is a continuous function of the roots of p, there exists a monic polynomial $p_0(x) \in \mathbb{P}(d)$ minimizing ϕ over degree d monic polynomials with roots in this interval. Observe that we can pick such p_0 having two distinct roots: if $\phi(p_0) < 0$ this is necessary by Claim C.3 and Equation (65), and if $\phi(p_0) = 0$ we can choose $p_0 = p$. By Corollary C.10, there exist $\widetilde{p}(x) \in \mathbb{P}(d)$ and $\widehat{p}(x) \in \mathbb{P}(d-1)$ such that $p_0(x) = \widehat{p}(x) + \widetilde{p}(x)$,

$$\mathsf{maxroot}(V_w p_0) = \mathsf{maxroot}(V_w \widetilde{p}) = \mathsf{maxroot}(V_w \widehat{p}), \tag{70}$$

and $\mathsf{maxroot}(\widetilde{p}) \leq \mathsf{maxroot}(p_0)$. By Lemma C.6 and by the bilinearity of the $\boxtimes$ operation,

$$\mathsf{maxroot}(V_w(p_0 \boxtimes q)) \leq \max\left\{\mathsf{maxroot}(V_w(\widehat{p} \boxtimes q)), \mathsf{maxroot}(V_w(\widetilde{p} \boxtimes q))\right\}, \tag{71}$$

with equality only if all three are equal. We note that

$$\begin{aligned}
\mathsf{maxroot}(V_w(\widehat{p} \boxtimes_d q)) &= \mathsf{maxroot}(V_w(\widehat{p} \boxtimes_{d-1} ((xD-d)q))) \quad \text{by Lemma C.7} \\
&\le \frac{w}{w+1} \mathsf{maxroot}(V_w \widehat{p}) \mathsf{maxroot}(V_w((xD-d)q)) \quad * \text{ induction hypothesis} \\
&\le \frac{w}{w+1} \mathsf{maxroot}(V_w \widehat{p}) \mathsf{maxroot}(V_w q)) \quad ** \text{ by Lemma C.8} \\
&= \frac{w}{w+1} \mathsf{maxroot}(V_w p_0) \mathsf{maxroot}(V_w q)) \quad \text{by Equation (70)} \\
&\le \mathsf{maxroot}(V_w(p_0 \boxtimes_d q)) \quad \text{since } \phi(p_0) \le 0.
\end{aligned}$$

As $*$ is a strict inequality in the case $d \ge 3$ (as both $\widehat{p}$ and $(xD-d)q$ have at least two distinct roots) and $**$ is strict in the case $d = 2$, we have by Equation (71) that $\mathsf{maxroot}(V_w(p_0 \boxtimes q)) < \mathsf{maxroot}(V_w(\widetilde{p} \boxtimes q))$, which implies $\phi(\widetilde{p}) < \phi(p_0)$, contradicting the minimality of p_0. □

ECCC ISSN 1433-8092
https://eccc.weizmann.ac.il