

# Large Low-Diameter Graphs are Good Expanders

Michael Dinitz<sup>1</sup>

Dept. of Computer Science, Johns Hopkins University, Baltimore, US  
mdinitz@cs.jhu.edu

Michael Schapira

School of Computer Science and Engineering, The Hebrew University, Jerusalem, Israel  
schapiram@cs.huji.ac.il

Gal Shahaf

Dept. of Mathematics, The Hebrew University, Jerusalem, Israel  
gal.shahaf@mail.huji.ac.il

---

## Abstract

We revisit the classical question of the relationship between the diameter of a graph and its expansion properties. One direction is well understood: expander graphs exhibit essentially the lowest possible diameter. We focus on the reverse direction, showing that “sufficiently large” graphs of fixed diameter and degree must be “good” expanders. We prove this statement for various definitions of “sufficiently large” (multiplicative/additive factor from the largest possible size), for different forms of expansion (edge, vertex, and spectral expansion), and for both directed and undirected graphs. A recurring theme is that the lower the diameter of the graph and (more importantly) the larger its size, the better the expansion guarantees. Aside from inherent theoretical interest, our motivation stems from the domain of network design. Both low-diameter networks and expanders are prominent approaches to designing high-performance networks in parallel computing, HPC, datacenter networking, and beyond. Our results establish that these two approaches are, in fact, inextricably intertwined. We leave the reader with many intriguing questions for future research.

**2012 ACM Subject Classification** Mathematics of computing → Spectra of graphs, Mathematics of computing → Extremal graph theory, Networks → Network design principles, Networks → Network structure

**Keywords and phrases** Network design, Expander graphs, Spectral graph theory

**Digital Object Identifier** 10.4230/LIPIcs.ESA.2018.71

**Acknowledgements** We wish to thank Nati Linial, Alex Samorodnitsky, Elchanan Mossel, and Noga Alon for fruitful discussions.

## 1 Introduction

Both the diameter of a graph and its expansion capture the “connectedness” of the graph, albeit in two very different senses. The diameter, i.e., the maximal distance between a pair of vertices, provides an upper bound on the length of shortest paths in the graph, whereas expansion measures the minimal ratio between a subset of vertices and its boundary. We revisit the classical question of relating these two traits. One direction is well known: good expansion implies a low diameter. Specifically, the diameter of a graph with good expansion

---

<sup>1</sup> [Supported in part by NSF awards 1464239 and 1535887]



is  $O(\log n)$  (see, e.g., [26]), which is asymptotically the lowest possible. We focus on the opposite, and largely unexplored, direction.

In general, low diameter does not guarantee good expansion. Consider, e.g., a graph on  $n$  vertices that is a disjoint union of two cliques, each of size  $\frac{n}{2}$ . Removing one edge from each clique and connecting the cliques via two “bridges” results in a  $(\frac{n}{2} - 1)$ -regular graph of diameter 3 with very low expansion (which worsens as  $n \rightarrow \infty$ ). We observe, however, that this “bad” graph is significantly smaller than the largest  $(\frac{n}{2} - 1)$ -regular graph of diameter 3 (which is of size  $\Omega(n^3)$ ). Indeed, our investigation below reveals that, in contrast to the above, when the degree and the diameter are fixed and the size of the graph is “sufficiently large”, the graph must have “good” expansion. We formalize this statement for different notions of “large”, for different forms of expansion (edge, vertex, and spectral expansion), and for undirected/directed graphs. Our results are presented in Section 1.2, but informally, “sufficiently large” means that the size of the graph is close to the best-known upper bound on the size, in either a multiplicative or additive setting.

We formalize the above statements and discuss implications of our results for network design and beyond, including the unification of two competing approaches to datacenter network design.

## 1.1 The Degree/Diameter Problem

Before we can state our results, we must first define what we mean by a “large, low-diameter graph”. To start off: how large can a  $d$ -regular graph of diameter  $k$  (which we shall refer to as a “ $(d, k)$ -graph” henceforth) actually be? An upper bound on the size of such a graph is the classical *Moore Bound* [25], denoted by  $\mu_{d,k}$  (see Section 2 for a formal definition). Extensive research has been devoted to determining the existence of graphs whose sizes match this upper bound (a.k.a., Moore Graphs) or well-approximate it. This line of study, termed the “*degree/diameter problem*”, was initiated by Hoffman and Singleton [25]. See [41] for a detailed survey of results in this active field of research.

Graphs whose sizes exactly match the Moore Bound, referred to as “Moore Graphs” henceforth, only exist for very few values of  $d$  and  $k$  [25, 4]. Consequently, various constructions for generating graphs whose sizes come “close” to the Moore Bound, which we call “approximate-Moore Graphs”, have been devised.

Specifically, constant multiplicative approximations (MMS-graphs [39]) and constant additive approximations (e.g., polarity graphs [20, 12]) have been devised for the case of diameter  $k = 2$ . Delorme [41, 16, 17] constructed infinite series of  $(d, k)$ -graphs whose sizes arbitrarily approach the Moore Bound for diameters  $k = 3$  and  $k = 5$ . Graph constructions whose sizes approximate the Moore Bound within non-constant multiplicative factors exist for arbitrary values of  $k$  (examples include, e.g., *de Bruijn* [15] and *Canale-Gomez* [13] graphs for the undirected case, and *Alegre* and [41] *Kautz* [19] digraphs for the directed analogue of the problem). While constructing approximate-Moore Graphs whose sizes arbitrarily approach the Moore Bound for arbitrary values of  $k$  remains an important and widely studied open question, such graphs are believed to exist for sufficiently large  $d$  and  $k$ , as conjectured, e.g., by Bollobás in [10].

Our investigation of the relation between diameter and expansion also uses the Moore Bound as a benchmark and compares the size of  $(d, k)$ -graphs to  $\mu_{d,k}$ . We consider both multiplicative and additive approximations to the Moore Bound. Our results establish that good solutions to the degree-diameter problem *must* be good expanders, establishing a novel link between two prominent and classical lines of research. In addition, our results yield new expansion bounds for all of the classical constructions of low-diameter graphs discussed above.

■ **Table 1** Summary of Results Relating Expansion and Diameter.

	Size	Expansion guarantees
$(d, k)$ -graph	$n \geq \mu_{d,k} - O(d^{k/2})$	$\lambda(G) = O(\sqrt{d})$
	$n \geq (1 - \varepsilon)\mu_{d,k}$	$\lambda(G) = O(\varepsilon^{1/k})d$
	$n = \alpha \cdot \mu_{d,k}$	$h_e(G) \geq \frac{\alpha d}{2k} \cdot \left(1 - \frac{1}{(d-1)^k}\right)$ $\phi_V(G) \geq \frac{\alpha}{2(k-1)+\alpha}$
$k = 2$	$n$	$\lambda(G) \leq \frac{1+\sqrt{1+4(d^2+d-n)}}{2}$
	$n = \alpha \cdot d^2$	$h_e(G) \geq \frac{2d+1-\sqrt{4(1-\alpha)d^2+4d+1}}{4}$ $\phi_V(G) \geq \frac{2\alpha}{2\alpha+1}$
$k = 3$	$n = \alpha \cdot d^3$	$\phi_V(G) \geq \frac{\alpha}{\alpha+1}$
$(d, k)$ -digraph	$n = \alpha \cdot \widetilde{\mu}_{d,k}$	$h_e(G) \geq \frac{\alpha}{2k} \left(d - \frac{1}{d^k}\right)$
		$\phi_V(G) \geq \frac{\alpha \cdot d}{2(d+1)(k-1)+\alpha \cdot d}$

## 1.2 Our Results

Our results relating the size of a  $(d, k)$ -graph to its expansion are summarized in Table 1, with  $\lambda(G)$ ,  $h_e(G)$ , and  $\phi_V(G)$  denoting spectral expansion, edge expansion, and vertex expansion, respectively (formal definitions can be found in Section 2).  $\widetilde{\mu}_{d,k}$  is the analogue of the Moore Bound for directed graphs.

We begin in Section 3 with our main results, which provide bounds on the spectral expansion of large  $(d, k)$ -graphs. In particular, our results establish that if the size of a graph is very close *additively* to the Moore bound, then the graph is essentially an optimal expander. In addition, if the graph has size that is close *multiplicatively* to the Moore bound, the spectral expansion might no longer be optimal, but is still very good.

We next turn our attention to combinatorial notions of expansion: edge expansion and vertex expansion. We provide (in Section 4) guarantees on both the edge and the vertex expansion of  $(d, k)$ -graphs in terms of their multiplicative distance from the Moore Bound. Our analysis leverages careful counting arguments to bound the ratio between the cardinality of a set of vertices and the size of its boundary. We also prove, through more refined analyses, improved results for diameters 2 and 3.

The key technical insight underlying our results for spectral expansion is a novel link, which we believe is of independent interest, between the nontrivial eigenvalues of a graph's adjacency matrix and the distance of the graph from the Moore Bound. Specifically, the proofs of our results for spectral expansion rely on the analysis of *non-backtracking paths* in the graph. A path is said to be non-backtracking if it does not traverse an edge back and forth consecutively. We prove that the matrix that corresponds to all non-backtracking paths of length at most  $k$  must consist of strictly positive entries and shares all eigenvectors of the adjacency matrix  $A$ . We establish the above algebraic relation between the two matrices by employing the *Geronimus Polynomials* [9, 44], a well-known class of orthogonal polynomials, as operators acting on the adjacency matrix. Given the spectrum of  $A$ , an asymptotic estimation of the polynomials' coefficients allows us to bound the spectrum of the non-backtracking paths matrix. We then subtract from the latter the all-ones matrix and use the leading eigenvalue of the remaining matrix (which can be computed directly) to bound the nontrivial eigenvalues of the adjacency matrix  $A$ .

■ **Table 2** Implications for Known Constructions of Low-Diameter Graphs.

Construction	Spectral expansion	Edge expansion	Vertex expansion
<i>de Bruijn</i> ( $k = 2$ ) [15]	-	-	$\frac{1}{3}$
<i>Canale-Gomez</i> [13]	-	$\frac{d}{2k \cdot 1.57^k} \left(1 - \frac{1}{(d-1)^k}\right)$	$\frac{1.57^{-k}}{2(k-1) + 1.57^{-k}}$
<i>Alegre digraph</i> [41]	-	$\frac{25 \cdot 2^k}{32k \cdot d^k} \left(d - \frac{1}{d^k}\right)$	$\frac{\left(\frac{2}{d}\right)^k \cdot \frac{25d}{16}}{2(d+1)(k-1) + \left(\frac{2}{d}\right)^k \cdot \frac{25d}{16}}$
<i>Kautz digraphs</i> [19]	-	$\frac{1}{2k} \left(d - \frac{1}{d^k}\right)$	$\frac{d}{2(d+1)(k-1) + d}$
<i>Polarity graph</i> [20, 12]	$\lambda \leq \frac{1 + \sqrt{1 + 8(d-1)}}{2}$	$\frac{2d+1 - \sqrt{4d+1}}{4}$	$\frac{2}{3}$
<i>MMS-graphs</i> [39]	$\lambda \leq \frac{1 + \frac{1}{3}\sqrt{d^2 + d + 7}}{2}$	$\frac{2d+1 - \sqrt{\frac{4}{9}d^2 + 4d + 1}}{4}$	$\frac{16}{25}$

Our technique should be contrasted with employing Hashimoto’s non-backtracking operator [24] to reason about non-backtracking paths in a graph (e.g., in the context of localization and centrality [38], clustering [33], mixing time acceleration [1], and percolation [27] in networks). In our context, applying Hashimoto’s operator involves reasoning about intricate relations between the spectra of the adjacency matrix  $A$  and another matrix, called the “non-backtracking matrix” (via the Ihara-Bass formula [11]). Geronimus Polynomials, a subfamily of the more renowned Chebyshev polynomials, allow for much simpler analysis. We believe that our techniques, and the (yet to be explored) connections between Geronimus Polynomials and Hashimoto’s operator, are of independent interest and may find wider applicability.

Importantly, our research diverges from the main vein of prior research on expanders. Expanders are commonly viewed as highly-connected sparse graphs. Indeed, the bulk of literature on this topic assumes that the degree of these graphs is essentially constant with respect to the size of the graph (i.e.,  $d \ll n$ ). In contrast, the size of a “large”  $(d, k)$ -graph is  $O(d^k)$ .

Aside from inherent theoretical interest, our motivation stems from the domain of network design. Low-diameter networks have been widely studied in the context of high-performance-computing (HPC) architectures (see, e.g., [29, 8, 2, 30]), parallel computing [34], and the design of fault-tolerant networks [6, 7, 8, 21, 42]. Of special interest in this literature are large networks of very low diameters (e.g., 2 or 3), as short path lengths translate to low latency in data delivery and also to low packet-queueing delays and power consumption (due to having few intermediate network devices en route to traffic destinations [20, 12, 39, 8, 29, 31]). Similarly to low-degree networks, expanders have been shown to induce high performance in a broad spectrum of network design contexts.

Recently, the focus on either the diameter or the expansion of a network topology gave rise to two competing approaches for datacenter architecture design [46, 43, 37, 32, 8, 29, 31]. Specifically, an important line of research in datacenter design (see, e.g., [43, 37, 8, 23]) relies on (either implicitly or explicitly) utilizing graphs whose sizes are as large as possible for a given diameter and degree as datacenter network topologies<sup>2</sup>. A different strand of research investigates how utilizing expander graphs as datacenter network topologies can be turned into an operational reality [46, 18, 28].

<sup>2</sup> The authors of [37], for instance, write that “Intuitively, the best known degree-diameter topologies should support a large number of servers with high network bandwidth and low cost (small degree)... Thus, we propose the best-known degree-diameter graphs as a benchmark for comparison.”

Our results show that these two approaches are, in fact, inextricably intertwined; not only do expanders exhibit low (in fact, near-optimal) diameters [26], but constructing large low-diameter datacenter networks effectively translates to constructing good expanders. Thus these two approaches to designing datacenter networks can essentially be regarded as one: the search for extremely strong expanders. Our results provide new expansion guarantees for a number of well-studied low-diameter networks, including MMS graphs [39] (proposed for the context high-performance computing and datacenters, see *Slim Fly* [8]), polarity graphs [20, 12], Canale-Gomez graphs [13], and more. See summary of the implications of our results for different graph constructions in Table 2. Our results for spectral expansion essentially match previously established results, thus generalizing and unifying prior construction-specific bounds.

Beyond the implications for network design, the study of low-diameter networks also pertains to other areas such as feedback registers [22, 35] and decoders [14].

## 2 Preliminaries

We provide below a brief exposition of graph expansion and the Moore Bound. We refer the reader to [26] and [41] for detailed expositions of these topics.

Let  $G = (V, E)$  be an undirected graph of size  $|V| = n$ .  $G$  is said to be  $d$ -regular if each of its vertices is of degree  $d$ , and of *diameter*  $k$  if the maximum distance between any two vertices in the graph is  $k$ .  $d$ -regular graphs of diameter  $k$  are denoted throughout the paper as  $(d, k)$ -graphs.

The combinatorial expansion of the graph reflects an isoperimetric view and is the minimal ratio between the boundary  $\partial S$  of a set  $S$  and its cardinality. Different interpretations of  $\partial S$  give rise to different notions of expansion.

The *edge expansion* of  $G$  is

$$h_e(G) := \min_{|S| \leq \frac{n}{2}} \frac{|e(S, S^c)|}{|S|}$$

where  $e(S, S^c) := \{(u, v) \in E \mid u \in S, v \in S^c\}$ .

The *vertex expansion* of  $G$  is

$$\phi_V(G) = \min_{0 < |S| \leq \frac{n}{2}} \frac{|N(S)|}{|S|}.$$

where  $N(S) := \{v \in S^c \mid \exists u \in S \text{ s.t. } (u, v) \in E\}$ .<sup>3</sup>

We next define the algebraic (spectral) notion of expansion. Let  $A$  be the adjacency matrix of the graph. Since  $A$  is symmetric it is diagonalizable with respect to an orthonormal basis, and the corresponding eigenvalues are real, and so can be ordered as follows:

$$\lambda_1 \geq \lambda_2 \geq \dots \geq \lambda_n.$$

The first eigenvalue of a  $d$ -regular graph satisfies  $\lambda_1 = d$  and has the all-ones vector  $\mathbf{1}_n$  as the associated eigenvector. Let  $\lambda(G) := \max\{|\lambda_2|, |\lambda_n|\}$ . A graph  $G$  is said to be an expander if  $\lambda(G)$  is bounded away from  $d$  by some constant [3]. The *algebraic expansion* (or spectral expansion) is then defined as  $d - \lambda(G)$ , termed the *spectral gap*<sup>4</sup>. The larger the gap, the better the expansion.

<sup>3</sup> The definitions of edge and vertex expansion admit several variants, based on either the size of the cut or the type of the boundary (see [26] for examples). While we adopt the most common of those, our results can be stated w.r.t. other variants as well.

<sup>4</sup> We use the two-sided notion of spectral gap throughout the paper, as opposed to  $d - \lambda_2$ .

## 71:6 Large Low-Diameter Graphs are Good Expanders

How large can a  $(d, k)$ -graph be? A straightforward upper bound is obtained by summation of the vertices according to their distance from a fixed vertex  $v_0 \in V$ . Let  $m_j$  denote the number of vertices at distance  $j$  from  $v_0$ . Note that  $m_0 = 1$  and  $m_1 = d$ . As vertices at distance  $j \geq 2$  must be adjacent to some vertex at distance  $j - 1$ , we have that  $m_j \leq (d - 1)m_{j-1}$ . A simple induction implies that  $m_j \leq d(d - 1)^{j-1}$ . Now since the diameter is  $k$ , all vertices have distance at most  $k$  from  $v_0$ , and hence  $n \leq 1 + d + d(d - 1) + d(d - 1)^2 + \dots + d(d - 1)^{k-1}$ . We denote this expression, known as the *Moore Bound* of the graph, by

$$\mu_{d,k} := 1 + d \sum_{i=0}^{k-1} (d-1)^i = \begin{cases} 2k + 1 & \text{if } d = 2 \\ 1 + d \cdot \frac{(d-1)^k - 1}{d-2} & \text{if } d > 2 \end{cases}$$

### 3 Diameter vs. Algebraic Expansion

We establish below a relationship between the nontrivial eigenvalues of  $A$  and the distance of the graph from the Moore Bound. This relationship will enable us to prove a variety of bounds on the algebraic expansion of approximate-Moore graphs. This novel link relies on the following class of orthogonal polynomials: let  $P_0(x) = 1$ ,  $P_1(x) = x$ ,  $P_2(x) = x^2 - d$ , and for every  $t > 2$  define  $P_t(x)$  by the recurrence relation

$$P_t(x) = xP_{t-1}(x) - (d-1)P_{t-2}(x).$$

The significance of this class of polynomials, termed the “*Geronimus Polynomials*” [9, 44], is reflected in the main technical theorem of this section:

► **Theorem 1.** *Let  $G$  be a  $(d, k)$ -graph of size  $n$ . Then, every nontrivial eigenvalue  $\lambda < d$  of  $G$  satisfies*

$$\left| \sum_{t=0}^k P_t(\lambda) \right| \leq \mu_{d,k} - n$$

Before delving into the proof of Theorem 1, we discuss some of its implications. Theorem 1 can be applied to provide meaningful guarantees regarding the spectral expansion of low-diameter graphs whose sizes approach the Moore Bound. Constructing large graphs of very low diameter, e.g.,  $k = 2, 3$ , has received much attention from both a theoretical perspective (see, e.g., [20, 12, 39]) and a practical perspective (see, e.g., [8, 29, 31]). An immediate implication of Theorem 1 is the following:

► **Theorem 2.** *Let  $G$  be a  $d$ -regular graph of diameter  $k = 2$  and size  $n$ , then*

$$\lambda(G) \leq \frac{1 + \sqrt{1 + 4(d^2 + d - n)}}{2}.$$

**Proof.** Applying the Geronimus Polynomials  $P_t(\lambda)$  for  $0 \leq t \leq 2$  in Theorem 1 yields

$$|1 + \lambda + (\lambda^2 - d)| \leq \mu_{d,2} - n = d^2 + 1 - n.$$

The result follows from solving the quadratic inequality. ◀

This theorem immediately bounds the algebraic expansion of polarity graphs [20, 12] and MMS graphs [39] claimed in Table 2, as both of these classes of graphs have diameter 2. What about graphs of diameter  $k > 2$ ? A more careful analysis of the Geronimus polynomials

for larger values of  $k$  allows us to use Theorem 1 to prove two different expansion bounds. The first is an extremely strong expansion bound but requires the size of the graph to be additively close to the Moore bound, whereas the second allows a small multiplicative gap between the size and the Moore bound but establishes a weaker expansion guarantee.

► **Theorem 3.** *Let  $G$  be a  $(d, k)$ -graph of size  $n \geq \mu_{d,k} - O(d^{k/2})$ , for some constant  $k > 0$ . Then  $\lambda(G) = O(\sqrt{d})$ .*

Since any  $d$ -regular graph must satisfy  $\lambda(G) \geq \sqrt{d}$  (see [26] for details), Theorem 3 implies that an additive approximation of  $O(d^{k/2})$  of the Moore Bound implies essentially optimal spectral properties.

► **Theorem 4.** *Let  $G$  be a  $(d, k)$ -graph of size  $n \geq (1 - \varepsilon)\mu_{d,k}$ , for some constant  $k > 0$ . Then  $\lambda(G) \leq O(\varepsilon^{1/k}) \cdot d$ .*

Delorme [41, 16, 17] proved the existence of an infinite series of  $(d, k)$ -graphs whose sizes arbitrarily approach the Moore Bound for diameters  $k = 3$  and  $k = 5$ . Specifically, Delorme proved that  $\liminf_{d \rightarrow \infty} \frac{n_{d,k}}{d^k} = 1$ , for  $k = 3, 5$ , where  $n_{d,k}$  is the largest possible size of a  $(d, k)$ -graph. This means that for  $k = 3, 5$ , for any constant  $\varepsilon > 0$  there is some value  $d'$  such that for all  $d \geq d'$  there is a  $(d, k)$ -graph with at least  $(1 - \varepsilon)\mu_{d,k}$  vertices. Hence, Theorem 4 implies that these graphs are good expanders.

Bollobás conjectured that  $(d, k)$ -graphs of size  $n \geq (1 - \varepsilon)d^k$  *always exist* for sufficiently large  $d$  and  $k$  [10]. Delorme's results may be perceived as supporting this conjecture. We point out that proving Bollobás' conjecture (or even extending Delorme's results to other specific values of  $k$  and  $d$ ) would immediately imply, by Theorem 4, similar expansion guarantees. The remainder of this section is devoted to the proofs of Theorems 1, 3, and 4.

### 3.1 Bounding the nontrivial eigenvalues (proof of Theorem 1)

Our high-level approach to proving Theorem 1 is the following: We aim to bound  $\lambda(G)$ , the second-largest eigenvalue (in absolute value) of the adjacency matrix  $A$ . We instead consider a different matrix  $M$ , obtained by employing the *Geronimus Polynomials* as operators over  $A$ . The combinatorial properties of this class of polynomials allow us to show that  $M\mathbf{1}_n = (\mu_{d,k} - n)\mathbf{1}_n$ . Applying the Perron-Frobenius Theorem asserts that this eigenvalue serves as a bound over the entire spectrum of  $M$ . We then utilize the algebraic relation between both matrices: Namely, we bound  $A$ 's nontrivial spectrum, using the fact that  $M$  shares the same eigenvectors as  $A$ , and that its eigenvalues may be derived from those of  $A$  via an operation of the Geronimus Polynomials. This will then imply Theorem 1.

We begin with the known solution to the recurrence, formulated via a trigonometric expression that holds for all  $t > 0$  [45]:

$$P_t(2\sqrt{d-1}\cos\theta) = (d-1)^{t/2-1} \frac{(d-1)\sin((t+1)\theta) - \sin((t-1)\theta)}{\sin\theta} \quad (1)$$

One can easily check that this identity applies for  $t = 1, 2$  and verify that the recurrence relation holds for  $t > 2$ . All roots of  $P_t$  are real and lie in the interval  $[-2\sqrt{d-1}, 2\sqrt{d-1}]$  [5, 36].

Our framework applies the Geronimus Polynomials as operators over the adjacency matrix  $A$ . This method has several advantages: Algebraically, since  $P_t(A)$  is a linear combination of powers of  $A$ , each eigenvector  $v$  of  $A$  is an eigenvector of  $P_t(A)$  as well. Thus, the spectrum of  $P_t(A)$  is given by  $\text{spec}[P_t(A)] = \{P_t(\lambda) \mid \lambda \text{ is an eigenvalue of } A\}$ . Viewed from

a combinatorial perspective, this operation allows us to dismiss backtracking paths from consideration. By *backtracking*, we refer to paths that traverse an edge in both directions consecutively. Note that a non-backtracking path need not be simple (a nontrivial cycle is a typical example of a non-backtracking yet non-simple path). The following claim states this observation formally. The proof is straightforward and is included for completeness.

► **Claim 5.** *Let  $A$  be the adjacency matrix of a  $d$ -regular graph  $G$ . Then,  $P_t(A)$  is the  $n \times n$  matrix in which the  $(u, v)$ 'th entry equals the number of non-backtracking paths of length exactly  $t$  between  $u$  and  $v$ .*

**Proof.** We use induction on  $t$ . Note that  $P_0(A) = I$ ,  $P_1(A) = A$  and  $P_2(A) = A^2 - dI$  satisfy the claim. For the induction step, suppose that the claim holds for all Geronimus Polynomials of order strictly less than  $t$ . Consider the term  $A \cdot P_{t-1}(A)$ , which corresponds to paths of length  $t$  such that the first  $t - 1$  hops on the path are non-backtracking. Note that reducible paths in this term are those paths that can only be reduced by eliminating their last two arcs and so there must be exactly  $(d - 1)P_{t-2}(A)$  of them. Being the difference between those quantities, it follows that  $P_t(A) = A \cdot P_{t-1}(A) - (d - 1)P_{t-2}(A)$  corresponds to the non-backtracking paths. ◀

As a corollary, the entries of  $P_t(A)$  are non-negative for all  $t \geq 0$ . In addition, as  $d(d - 1)^{t-1}$  is the number of non-backtracking paths of length  $t > 0$  starting from every  $v \in G$ , this quantity equals the sum of entries in every row of  $P_t(A)$ . Hence, Claim 5 implies that  $P_t(A)\mathbf{1}_n = d(d - 1)^{t-1}\mathbf{1}_n$ .

Summing over the indices  $0 \leq t \leq k$ , yields

$$\sum_{t=0}^k P_t(A)\mathbf{1}_n = \left(1 + \sum_{t=1}^k d(d - 1)^{t-1}\right)\mathbf{1}_n = \mu_{d,k} \cdot \mathbf{1}_n. \tag{2}$$

We are now ready to prove Theorem 1:

**Proof. (of Theorem 1)** Given Claim 5, the sum of matrices  $\sum_{t=0}^k P_t(A)$  corresponds to all non-backtracking paths of length at most  $k$ . Since  $G$  is of diameter  $k$ , this sum of matrices must consist of strictly positive entries, and can thus be represented as  $\sum_{t=0}^k P_t(A) = J + M$ , where  $J$  is the all ones matrix and  $M$  is non-negative. We now have  $M\mathbf{1}_n = \left(\sum_{t=0}^k P_t(A) - J\right)\mathbf{1}_n = (\mu_{d,k} - n)\mathbf{1}_n$ , where the second equality is due to (2).

Recall that  $A$  is symmetric and thus diagonalizable w.r.t. an orthogonal basis. Therefore, any eigenvector  $v \neq \mathbf{1}_n$  must lie in  $(\text{span}\{\mathbf{1}_n\})^\perp$ . Since  $J\mathbf{1}_n = n\mathbf{1}_n$  and  $\text{rank}(J) = 1$ , it follows that  $Jv = 0$ . Hence,  $Mv = \left(\sum_{t=0}^k P_t(A) - J\right)v = \sum_{t=0}^k P_t(A)v = \sum_{t=0}^k P_t(\lambda)v$ .

This implies in particular that

$$\text{spec}(M) = \left\{ \sum_{t=0}^k P_t(\lambda) \mid \lambda \text{ is a nontrivial eigenvalue of } A \right\} \cup \{\mu_{d,k} - n\}.$$

We now apply the Perron-Frobenius Theorem (see [40]), which states that a non-negative matrix admits a non-negative eigenvector with a non-negative eigenvalue that is larger or equal, in absolute value, to all other eigenvalues. Now, since  $\mathbf{1}_n$  is the only non-negative eigenvector of  $M$ , we conclude that  $\mu_{d,k} - n$  is the leading eigenvalue of  $M$  and the claim follows. ◀



### 3.2 Proof of Theorem 3

Our proof of Theorem 3 utilizes a careful asymptotic estimation of the Geronimus Polynomials' coefficients. When  $\lambda(G)$  is of order larger than  $\sqrt{d}$ , our analysis asserts that  $\left| \sum_{t=0}^k P_t(\lambda) \right|$  must be larger than  $O(d^{k/2})$  for some nontrivial eigenvalue  $\lambda$  of  $G$ , thus resulting in a contradiction to Theorem 1.

For our purposes, it will be beneficial to use the representation  $P_t(x) = \sum_{i=0}^t a_{t,i} x^i$ , where  $a_{t,i}$  is the  $i$ 'th coefficient of the  $t$ 'th Geronimus Polynomial. We note the following: (i)  $P_t$  is either odd or even<sup>5</sup> for all  $t > 0$ , and the parity of  $P_t$  equals the parity of  $t$ . This can be shown either by induction using the recurrence relation, or straightforward from the solution (1); (ii) A comparison of the leading coefficients in the recurrence implies that  $a_{t,t} = a_{t-1,t-1}$ . Applying the boundary conditions ( $a_{1,1} = a_{0,0} = 1$ ) yields  $a_{t,t} = 1$  for all  $t > 0$ ; (iii) Setting  $\theta = \frac{\pi}{2}$  in (1) yields  $a_{t,0} = d(d-1)^{t/2-1}(-1)^{t/2}$  whenever  $t$  is even.

The following easy-to-prove claim provides us with asymptotic estimates for the rest of the coefficients. Note that the  $\Theta(\cdot)$  notation is hiding factors of  $t$  (we will only use this claim only where  $t$  is constant).

► **Claim 6.** Let  $P_t(x) = \sum_{i=0}^t a_{t,i} x^i$  denote the Geronimus polynomial of order  $t$ , then

$$a_{t,i} = \begin{cases} (-1)^{\frac{t-i}{2}} \Theta\left(d^{\frac{t-i}{2}}\right) & \text{if } (t-i) \text{ is even} \\ 0 & \text{if } (t-i) \text{ is odd} \end{cases}$$

for all  $0 \leq i \leq t$ .

This immediately gives the next corollary, which is just a slightly easier to use formulation of  $P_t(x)$ .

► **Corollary 7.** The Geronimus Polynomial of order  $t$  can be written as

$$P_t(x) = \sum_{i=0}^{\lfloor \frac{t}{2} \rfloor} (-1)^i \cdot \Theta(d^i) \cdot x^{t-2i}.$$

We are now ready to apply this machinery. The following lemma bounds the value of these polynomials on values which are “small”.

► **Claim 8.** Let  $\frac{1}{2} < \alpha \leq 1$ , and let  $|\lambda| = \Theta(d^\alpha)$ . Then  $|P_t(\lambda)| = \Theta(d^{t\alpha})$ .

**Proof.** We use induction on  $t$ . For  $t = 0$  we have that  $P_0(\lambda) = 1 = \Theta(d^{0\alpha})$ , and for  $t = 1$  we have that  $|P_1(\lambda)| = |\lambda| = |\Theta(d^{1\alpha})|$ . Assume that the claim holds for the Geronimus Polynomials of order less than  $t$ . Using Corollary 7, we now have

$$P_t(\lambda) = \sum_{i=0}^{\lfloor \frac{t}{2} \rfloor} (-1)^i \Theta(d^i) \lambda^{t-2i} = \sum_{i=0}^{\lfloor \frac{t}{2} \rfloor} (-1)^i \Theta(d^i) \Theta(d^{\alpha(t-2i)}) = \sum_{i=0}^{\lfloor \frac{t}{2} \rfloor} (-1)^i \cdot \Theta(d^{t\alpha+i(1-2\alpha)}).$$

Whenever  $\alpha > \frac{1}{2}$ , the absolute value of this equals  $\Theta(d^{t\alpha})$  as claimed. ◀

**Proof. (of Theorem 3)** Suppose that  $A$  obtains an eigenvalue  $\lambda = \Theta(d^\alpha)$  for some  $\alpha > \frac{1}{2}$ . Then, applying Claim 8, we have:

$$\left| \sum_{t=0}^k P_t(\lambda) \right| = \left| \sum_{t=0}^k \Theta(d^{t\alpha}) v \right| = \Theta(d^{k\alpha}) v$$

<sup>5</sup> A polynomial  $q(x)$  is said to be even if  $q(x) = q(-x)$  and odd if  $q(-x) = -q(x)$ .

## 71:10 Large Low-Diameter Graphs are Good Expanders

This expression, however, is upper bounded by  $\mu_{d,k} - n$  (by Theorem 1), which is  $O(d^{k/2})$  by the assumption of the Theorem 3. We thus have

$$|\Theta(d^{k\alpha})| \leq \mu_{d,k} - n \leq O(d^{k/2})$$

and this is, of course, a contradiction to the assumption  $\alpha > \frac{1}{2}$ . We therefore conclude that  $\lambda = O(\sqrt{d})$ . ◀

### 3.3 Proof of Theorem 4

The proof of Theorem 4 relies on some of the ideas introduced in the proof of Theorem 3. Let  $\lambda$  be a nontrivial eigenvalue of  $G$ . We wish to show that  $|\lambda| \leq O(\varepsilon^{1/k})d$ . If  $|\lambda| \leq O(d^{2/3})$  then we are done. Suppose, then, that  $|\lambda| \geq \omega(d^{2/3})$ , and hence  $d = o(|\lambda|^{3/2})$ . Consider the sum  $|\sum_{t=0}^k P_t(\lambda)|$ . Corollary 7, and the discussion which showed that  $a_{t,t} = 1$ , imply that this sum is at least

$$\begin{aligned} \left| \sum_{t=0}^k P_t(\lambda) \right| &\geq |\lambda^k + \lambda^{k-1}| - \sum_{i=1}^{\lfloor k/2 \rfloor} (\Theta(d^i)|\lambda|^{k-2i} + \Theta(d^i)|\lambda|^{k-2i-1}) \\ &\geq |\lambda^k + \lambda^{k-1}| - \sum_{i=1}^{\lfloor k/2 \rfloor} \left( \Theta(|\lambda|^{k-(i/2)}) + \Theta(|\lambda|^{k-1-(i/2)}) \right) \geq \Theta(|\lambda^k|), \end{aligned}$$

where the second inequality follows from the assumption that  $|\lambda| \geq \omega(d^{2/3})$ .

When we plug this into Theorem 1, we get that  $\Theta(|\lambda^k|) \leq \mu_{d,k} - n \leq \varepsilon\mu_{d,k}$ . Since  $\mu_{d,k} \leq cd^k$  for some constant  $c$ , we get that  $c'|\lambda^k| \leq \varepsilon cd^k$  for some constants  $c$  and  $c'$ , and hence  $|\lambda| \leq \left(\frac{c}{c'}\right)^{1/k} \varepsilon^{1/k} d$ , proving the theorem. ◀

## 4 Diameter vs. Combinatorial Expansion

We present below our results for combinatorial expansion. We first point out that applying the Cheeger inequality [26] to our bounds on spectral expansion immediately implies bounds on combinatorial expansion. Specifically, the Cheeger inequality states that  $h_e(G) \geq \frac{d-\lambda_2}{2}$ . When combined with Theorems 3 and 4, this yields the following bounds.

► **Theorem 9.** *Let  $G$  be a  $(d, k)$  graph with  $n$  vertices, for some constant  $k > 0$ . If  $n \geq \mu_{d,k} - O(d^{k/2})$  then  $h_e(G) \geq \frac{d-O(\sqrt{d})}{2}$ .*

► **Theorem 10.** *Let  $G$  be a  $(d, k)$  graph with  $n$  vertices, for some constant  $k > 0$ . If  $n \geq (1 - \varepsilon)\mu_{d,k}$  then  $h_e(G) \geq \frac{(1-O(\varepsilon^{1/k}))d}{2}$ .*

Observe that, since clearly  $d/2$  is an upper bound on  $h_e(G)$ , both of these bounds imply very high expansion guarantees when  $n$  is very close to the Moore Bound. However, when this is not so, e.g., when  $n = \mu_{d,k}/k$ , neither bound yields nontrivial expansion guarantees.

To provide stronger expansion guarantees for graphs that do not come very close (additively/ multiplicatively) to the Moore Bound, we analyze combinatorial expansion directly. We next present our bounds for edge and vertex expansion in undirected and directed graphs. We discuss the implications of these expansion bounds for known  $(d, k)$ -graph constructions in Table 2.

Our main result of this section is the following:

► **Theorem 11.** Let  $G = (V, E)$  be a  $d$ -regular graph of size  $n$  and diameter  $k$ . If  $n = \alpha \cdot \mu_{d,k}$ , then

$$h_e(G) \geq \frac{\alpha d}{2k} \cdot \left(1 - \frac{1}{(d-1)^k}\right) \quad \text{and} \quad \phi_V(G) \geq \frac{\alpha}{2(k-1) + \alpha}.$$

Our proof of Theorem 11 utilizes a counting argument. As the graph has diameter  $k$ , each pair of vertices on opposite sides of a cut must be connected via a path of length at most  $k$  that traverses the boundary. However, there is an upper bound, induced by the degree and diameter of the graph, on the number of such paths that traverse a given edge/vertex. A careful examination of the implications of these two limitations provides us with a lower bound on the size of the boundary.

**Proof.** (first part of Theorem 11) Let  $(S, S^c)$  be a cut in the graph, and let  $|S| = s \leq \frac{n}{2}$ . As the diameter equals  $k$ , every pair of vertices that lie on both sides of the cut must be connected via a path of length at most  $k$ . We thus have  $s(n-s)$  such paths, each of which passes through some edge in the cut.

How many paths of length at most  $k$  include a given edge  $e \in E$ ? As  $G$  is  $d$ -regular, there are at most  $(d-1)^{l-1}$  paths of length  $l$  for which  $e$  is in the  $i$ 'th position in the path. It follows that no more than  $l \cdot (d-1)^{l-1}$  paths of length  $l$  use a specific edge, hence the number of paths of length at most  $k$  that utilize a fixed edge is upper bounded by

$$f_{d-1}(k) = \sum_{l=1}^k l \cdot (d-1)^{l-1}.$$

Let us find a simpler formulation of  $f_{d-1}(k)$ . Integrating yields

$$F_{d-1}(k) = \sum_{l=1}^k (d-1)^l = \frac{(d-1)^{k+1} - 1}{(d-1) - 1}.$$

Differentiating brings us back to

$$\begin{aligned} f_{d-1}(k) &= \frac{(k+1)(d-1)^k(d-2) - [(d-1)^{k+1} - 1]}{(d-2)^2} \\ &\leq \frac{(k+1)(d-1)^k(d-2) - (d-1)^k(d-2)}{(d-2)^2} \\ &= \frac{k(d-1)^k}{(d-2)} \end{aligned}$$

Now,  $s(n-s)$  paths use the cut, and every edge in the cut can be a part of at most  $f_{d-1}(k)$  paths. It follows that  $|e(S, S^c)| \geq \frac{s(n-s)}{f_{d-1}(k)}$  for every cut  $(S, S^c)$  in  $G$ . Hence, the cut that realizes  $h_e(G)$  satisfies

$$\begin{aligned} h_e(G) &= \frac{|e(S, S^c)|}{|S|} \geq \frac{s(n-s)(d-2)}{s \cdot k(d-1)^k} \geq \frac{n}{2} \cdot \frac{(d-2)}{k(d-1)^k} \\ &= \frac{\alpha d \cdot ((d-1)^k - 1)}{2(d-2)} \cdot \frac{(d-2)}{k(d-1)^k} \\ &= \frac{\alpha d}{2k} \cdot \left(1 - \frac{1}{(d-1)^k}\right). \end{aligned}$$

#### 4.1 Directed graphs.

We consider directed graphs next. We begin by introducing the relevant terminology and notation. We say that a directed graph (a.k.a. digraph)  $G = (V, E)$  is  $d$ -regular if both the out-degree and the in-degree of each vertex equals  $d$ . A cut in a digraph  $e(S, S^c) = \{(u, v) \in E \mid u \in S, v \in S^c\}$  is asymmetric, and consists of all edges directed from  $S$  to  $S^c$ . The diameter is still defined as the maximal distance between two vertices, and the corresponding Moore Bound is only slightly different (as there are potentially  $d^i$  vertices of distance  $i$  from a given vertex):  $\tilde{\mu}_{d,k} = \sum_{i=0}^k d^i = \frac{d^{k+1}-1}{d-1}$ .

The following result is the directed analogue of Theorem 11:

► **Theorem 12.** *Let  $G$  be a  $d$ -regular,  $k$ -diameter **directed** graph of size  $n = \alpha \cdot \tilde{\mu}_{d,k}$ , then*

$$h(G) \geq \frac{\alpha}{2k} \left( d - \frac{1}{d^k} \right) \quad \text{and} \quad \phi_V(G) \geq \frac{\alpha \cdot d}{2(d+1)(k-1) + \alpha \cdot d}.$$

Much research on constructing low-diameter graphs focuses on diameters 2 and 3 (see, e.g., [41, 20, 39]). Graphs of very low diameter are particularly important from a practical perspective [8, 29, 31]. The following theorems improve upon our results for the edge expansion and vertex expansion of  $(d, k)$ -graphs.

► **Theorem 13.** *Let  $G = (V, E)$  be an undirected  $(d, 2)$ -graph of size  $n = \alpha \cdot d^2$ . Then*

$$h_e(G) \geq \frac{2d+1 - \sqrt{4(1-\alpha)d^2 + 4d+1}}{4}.$$

► **Theorem 14.** *Let  $G = (V, E)$  be an undirected  $(d, 2)$ -graph of size  $n = \alpha \cdot d^2$ . Then  $\phi_V(G) \geq \frac{2\alpha}{2\alpha+1}$ .*

We can extend our analysis to graphs of diameter 3, yielding the following theorem.

► **Theorem 15.** *Let  $G = (V, E)$  be a  $(d, 3)$ -graph of size  $n = \alpha \cdot d^3$ , then  $\phi_V(G) \geq \frac{\alpha}{\alpha+1}$ .*

## 5 Conclusion and Open Questions

We revisited the classical question of relating the expansion and the diameter of graphs and showed that not only do good expanders exhibit low diameter but the converse is also, in some sense, true. We also discussed the implications of our results for constructions from the rich body literature on low-diameter graphs. We leave the reader with many interesting open questions, including: (1) **Tightening the gaps.** An obvious open question is improving upon our lower bounds and establishing upper bounds on the expansion of fixed-diameter graph constructions. (2) **Benchmarking against the optimal (largest possible)  $(d, k)$ -graph.** We used the Moore Bound as a benchmark. Another approach would be to compare against the size of the largest possible  $(d, k)$ -graph. (3) **Geronimus Polynomials vs. Hashimoto's non-backtracking operator.** The operation of the Geronimus Polynomials over the adjacency matrix of the graph offers a new perspective on its non-backtracking paths (as established in Lemma 5). This suggests a non-trivial relation between these polynomials and Hashimoto's non-backtracking operator, which we believe is of independent interest and may find wider applicability.

## References

- 1 Noga Alon, Itai Benjamini, Eyal Lubetzky, and Sasha Sodin. Non-backtracking random walks mix faster. *Communications in Contemporary Mathematics*, 9(04):585–603, 2007.
- 2 Baba Arimilli, Ravi Arimilli, Vicente Chung, Scott Clark, Wolfgang Denzel, Ben Drerup, Torsten Hoefler, Jody Joyner, Jerry Lewis, Jian Li, et al. The percs high-performance interconnect. In *High Performance Interconnects (HOTI), 2010 IEEE 18th Annual Symposium on*, pages 75–82. IEEE, 2010.
- 3 Sanjeev Arora and Boaz Barak. *Computational complexity: a modern approach*. Cambridge University Press, 2009.
- 4 Eiichi Bannai and Tatsuro Ito. Regular graphs with excess one. *Discrete Mathematics*, 37(2):147–158, 1981.
- 5 Yair Bartal, Nathan Linial, Manor Mendel, and Assaf Naor. On metric ramsey-type phenomena. In *Proceedings of the thirty-fifth annual ACM symposium on Theory of computing*, pages 463–472. ACM, 2003.
- 6 J-C Bermond, Nathalie Homobono, and Claudine Peyrat. Large fault-tolerant interconnection networks. *Graphs and Combinatorics*, 5(1):107–123, 1989.
- 7 Jean-Claude Bermond. De bruijn and kautz networks: a competitor for the hypercube? *Hypercube and distributed computers*, pages 279–293, 1989.
- 8 Maciej Besta and Torsten Hoefler. Slim fly: A cost effective low-diameter network topology. In *Proceedings of the International Conference for High Performance Computing, Networking, Storage and Analysis*, pages 348–359. IEEE Press, 2014.
- 9 Norman Biggs. *Algebraic graph theory*. Cambridge university press, 1993.
- 10 B Bollobás. *Extremal Graph Theory*. Dover Publications, 1978.
- 11 C Bordenave. A new proof of friedman’s second eigenvalue theorem and its extension to random lifts. *Preprint, available at <https://arxiv.org/abs/1502.04482>*, 2015.
- 12 William G Brown. On graphs that do not contain a thomsen graph. *Canad. Math. Bull.*, 9(2):1–2, 1966.
- 13 Eduardo A Canale and José Gómez. Asymptotically large  $(\delta, d)$ -graphs. *Discrete applied mathematics*, 152(1):89–108, 2005.
- 14 Oliver Collins, Sam Dolinar, Robert McEliece, and Fabrizio Pollara. A vlsi decomposition of the de bruijn graph. *Journal of the ACM (JACM)*, 39(4):931–948, 1992.
- 15 DG De Bruijn. A combinatorial problem, koninklijke nederlandse academie van wetenschappen et amsterdam. *Proceedings Qt the Section gt Sciences*, 49:27, 1946.
- 16 Charles Delorme. Grands graphes de degré et diametre donnés. *European Journal of Combinatorics*, 6(4):291–302, 1985.
- 17 Charles Delorme. Large bipartite graphs with given degree and diameter. *Journal of graph theory*, 9(3):325–334, 1985.
- 18 Michael Dinitz, Michael Schapira, and Asaf Valadarsky. Explicit expanding expanders. *Algorithmica*, pages 1–21, 2016. doi:10.1007/s00453-016-0269-x.
- 19 Bernard Elspas, William H Kautz, and James Turner. Theory of cellular logic networks and machines. Technical report, DTIC Document, 1968.
- 20 Paul Erdos, Alfréd Rényi, and VT Sós. On a problem in the theory of graphs. *Publ. Math. Inst. Hungar. Acad. Sci.*, 7:215–235, 1962.
- 21 A-H Esfahanian and S. Louis Hakimi. Fault-tolerant routing in de bruijn communication networks. *IEEE Transactions on Computers*, 100(9):777–788, 1985.
- 22 Harold Fredricksen. A survey of full length nonlinear shift register cycle algorithms. *SIAM review*, 24(2):195–221, 1982.
- 23 D. Guo, T. Chen, D. Li, M. Li, Y. Liu, and G. Chen. Expandable and cost-effective network structures for data centers using dual-port servers. *IEEE Transactions on Computers*, 62(7):1303–1317, July 2013. doi:10.1109/TC.2012.90.

- 24 Ki-ichiro Hashimoto. Zeta functions of finite graphs and representations of p-adic groups. *Automorphic forms and geometry of arithmetic varieties.*, pages 211–280, 1989.
- 25 Alan J Hoffman and Robert R Singleton. On moore graphs with diameters 2 and 3. *IBM Journal of Research and Development*, 4(5):497–504, 1960.
- 26 Shlomo Hoory, Nathan Linial, and Avi Wigderson. Expander graphs and their applications. *Bulletin of the American Mathematical Society*, 43(4):439–561, 2006.
- 27 Brian Karrer, Mark EJ Newman, and Lenka Zdeborová. Percolation on sparse networks. *Physical review letters*, 113(20):208702, 2014.
- 28 Simon Kassing, Asaf Valadarsky, Gal Shahaf, Michael Schapira, and Ankit Singla. Beyond fat-trees without antennae, mirrors, and disco-balls. In *Proceedings of the Conference of the ACM Special Interest Group on Data Communication*, pages 281–294. ACM, 2017.
- 29 John Kim, William J Dally, Steve Scott, and Dennis Abts. Technology-driven, highly-scalable dragonfly topology. In *ACM SIGARCH Computer Architecture News*, volume 36, pages 77–88. IEEE Computer Society, 2008.
- 30 John Kim, William J Dally, and Dennis Abts. Flattened butterfly: a cost-efficient topology for high-radix networks. In *ACM SIGARCH Computer Architecture News*, volume 35, pages 126–137. ACM, 2007.
- 31 John Kim, William J Dally, Steve Scott, and Dennis Abts. Cost-efficient dragonfly topology for large-scale systems. In *Optical Fiber Communication Conference*, page OTuI2. Optical Society of America, 2009.
- 32 Michihiro Koibuchi, Hiroki Matsutani, Hideharu Amano, D Frank Hsu, and Henri Casanova. A case for random shortcut topologies for hpc interconnects. In *Computer Architecture (ISCA), 2012 39th Annual International Symposium on*, pages 177–188. IEEE, 2012.
- 33 Florent Krzakala, Christopher Moore, Elchanan Mossel, Joe Neeman, Allan Sly, Lenka Zdeborová, and Pan Zhang. Spectral redemption in clustering sparse networks. *Proceedings of the National Academy of Sciences*, 110(52):20935–20940, 2013.
- 34 F Thomson Leighton. *Introduction to parallel algorithms and architectures: Arrays · trees · hypercubes*. Elsevier, 2014.
- 35 Abraham Lempel. On a homomorphism of the de bruijn graph and its applications to the design of feedback shift registers. *IEEE Transactions on Computers*, 100(12):1204–1209, 1970.
- 36 Nathan Linial, Avner Magen, and Assaf Naor. Girth and euclidean distortion. *Geometric & Functional Analysis GAFA*, 12(2):380–394, 2002.
- 37 Ankit Singla Chi-Yao Hong Lucian and Popa Brighten Godfrey. Jellyfish: Networking data centers randomly. *CoRR*, abs/1110.1687, 2011. URL: <http://arxiv.org/abs/1110.1687>.
- 38 Travis Martin, Xiao Zhang, and MEJ Newman. Localization and centrality in networks. *Physical Review E*, 90(5):052808, 2014.
- 39 Brendan D McKay, Mirka Miller, and Jozef Širáň. A note on large graphs of diameter two and given maximum degree. *Journal of Combinatorial Theory, Series B*, 74(1):110–118, 1998.
- 40 Carl D Meyer. *Matrix analysis and applied linear algebra*, volume 2. Siam, 2000.
- 41 Mirka Miller and Jozef Širáň. Moore graphs and beyond: A survey of the degree/diameter problem. *Electronic Journal of Combinatorics*, 61:1–63, 2005.
- 42 Dhiraj K Pradhan. Fault-tolerant multiprocessor link and bus network architectures. *IEEE Trans. Comput.:(United States)*, 100, 1985.
- 43 Ankit Singla, P Brighten Godfrey, and Alexandra Kolla. High throughput data center topology design. In *11th USENIX Symposium on Networked Systems Design and Implementation (NSDI 14)*, pages 29–41, 2014.
- 44 Patrick Solé. The second eigenvalue of regular graphs of given girth. *Journal of Combinatorial Theory, Series B*, 56(2):239–249, 1992.

- 45 Gabor Szeg. *Orthogonal polynomials*, volume 23. American Mathematical Soc., 1939.
- 46 Asaf Valadarsky, Gal Shahaf, Michael Dinitz, and Michael Schapira. Xpander: Towards optimal-performance datacenters. In *Proceedings of the 12th International on Conference on Emerging Networking EXperiments and Technologies*, CoNEXT '16, pages 205–219, New York, NY, USA, 2016. ACM. doi:10.1145/2999572.2999580.