



Ramanujan graphs and the spectral gap of supercomputing topologies

Sinan G. Aksoy¹ · Paul Bruillard¹ · Stephen J. Young¹ · Mark Raugas²

© This is a U.S. government work and not under copyright protection in the U.S.; foreign copyright protection may apply 2020

Abstract

Graph eigenvalues play a fundamental role in controlling structural properties which are critical considerations in the design of supercomputing interconnection networks, such as bisection bandwidth, diameter, and fault tolerance. This motivates considering graphs with optimal spectral expansion, called *Ramanujan graphs*, as potential candidates for interconnection networks. In this work, we explore this possibility by comparing Ramanujan graph properties against those of a wide swath of current and proposed supercomputing topologies. We derive analytic expressions for the spectral gap, bisection bandwidth, and diameter of these topologies, some of which were previously unknown. We find the spectral gap of existing topologies is well separated from the optimal achievable by Ramanujan topologies, suggesting the potential utility of adopting Ramanujan graphs as interconnection networks.

Keywords Ramanujan graphs · Expander graphs · Supercomputing topologies · Interconnection networks

✉ Sinan G. Aksoy
sinan.aksoy@pnnl.gov

Paul Bruillard
bruillardp@gmail.com

Stephen J. Young
stephen.young@pnnl.gov

Mark Raugas
mark.raugas@pnnl.gov

¹ Pacific Northwest National Laboratory, Richland, WA 99352, USA

² Pacific Northwest National Laboratory, Seattle, WA 98109, USA

1 Introduction

One of the significant challenges in the use of modern cluster-based supercomputers is efficiently, robustly, and quickly handling the necessary communication between nodes in the cluster. Both the current and next-generation supercomputer designs use highly structured network topologies, such as the low-dimensional torus, the flattened butterfly, or the dragonfly topology in order to have a straightforward routing scheme while attempting to mitigate the traffic congestion in high communication applications. However, “preliminary experiments on Edison, a Cray XC30 at NERSC, have shown that for communications-heavy applications, inter-job interference and thus network congestion remain an important factor” [12]. Indeed, recent research [55] further attests to the impact of network structure on performance metrics. In fact, even with a relatively low utilization (40–50%), communication patterns can cause an exponential explosion in latency [37]. As a consequence of the interaction between the structure of internode communication in various classes of algorithms and the underlying network topologies, certain supercomputers gain a reputation for being more or less suited to a certain class of problems.

In this regard, the evolution of supercomputing interconnection topologies stands in contrast to the surprising success of the “evolved” topology of the Internet. Specifically, despite having no global design, the Internet structure has unexpectedly [2] ended up as a robust, general purpose, and relatively low-latency system for its size. In the last few decades, a consensus has developed that the primary explanation for the good performance of the Internet topology is that the Internet topology belongs to a class of graphs known as *expanders*. That is, if a graph is a sufficiently high-quality expander, then there exist efficient, distributed, online, local, and low-congestion algorithms to route information among the vertices of the graph [7,21,26,29,38,64].

This view point leads naturally to considering optimal expanders, known as Ramanujan graphs, as potential supercomputing topologies. In this work, we explore the potential benefits of adopting Ramanujan graphs by conducting an analysis of current and proposed supercomputing topologies. The paper is organized as follows: In Sect. 2, we provide the necessary preliminaries on spectral graph theory, as well as survey results showing eigenvalues control a number of critical properties pertinent to interconnection design, such as bisection bandwidth, diameter, and fault tolerance. Second, in Sect. 3 we define the Ramanujan property of graphs and review explicit constructions of Ramanujan graphs. In Sect. 4, we survey variety of supercomputing topologies and derive analytic expressions for their spectral expansion, bisection bandwidth, and diameter. Across the topologies surveyed, we find some or all of these properties are well separated from those of Ramanujan topologies. Consequently, our results suggest transition to Ramanujan topologies may have the potential to significantly improve metrics for facility of communication.

2 Preliminaries

Before proceeding with our discussion of expanders and Ramanujan graphs, we first recall some relevant terminology and results from graph theory. A graph $G = (V, E)$




Fig. 1 How eigenvalues are obtained from a graph is represented by the solid lines. Spectral graph theory analyzes eigenvalues to deduce properties of the graph, represented by the dashed line

is a set of vertices V edges E , where each edge is an unordered pair of vertices. The number of edges incident to a vertex is called its *degree*; if every vertex has degree k , the graph is called k -regular. Spectral graph theory is the study of eigenvalues and eigenvectors of matrices associated with graphs. The *adjacency matrix* A of an n -vertex graph is an $n \times n$ matrix where

$$A_{ij} = \begin{cases} 1 & \text{if } \{i, j\} \in E \\ 0 & \text{otherwise} \end{cases} .$$

As A is symmetric, its eigenvalues are real, which we denote

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

For a connected graph, the largest eigenvalue $\lambda_1 = k$ if and only if the graph is k -regular; furthermore, if G is connected, $\lambda_1 - \lambda_2 > 0$, and the quantity $\lambda_1 - \lambda_2$ is referred to as the *spectral gap* of G .

Two other graph matrices whose spectra are often studied are the *Laplacian matrix* $L = D - A$ and *normalized Laplacian matrix* $\mathcal{L} = D^{-1/2} L D^{-1/2}$, where D denotes the diagonal matrix with the vertex degrees on the diagonal. Unlike the adjacency matrix, both of these matrices are necessarily positive semi-definite, and their spectra characterize a number of properties which are not captured by adjacency eigenvalues. Due to its intimate connection to random walks and stochastic processes on graphs, the normalized Laplacian matrix is perhaps the most appropriate matrix for characterizing expansion properties of graphs, particularly for irregular graphs. However, we note that if a graph G is k -regular (as is the case for a number of supercomputing topologies), then $\mathcal{L} = I - \frac{1}{k} A$, from which it is clear that the spectra of all three matrices are related by trivial shifts and scalings by k (and hence functionally the same). We denote the spectrum of the Laplacian matrix L by

$$0 = \rho_1 \leq \rho_2 \leq \dots \leq \rho_n ,$$

and that of the normalized Laplacian \mathcal{L} by

$$0 = \mu_1 \leq \mu_2 \leq \dots \leq \mu_n \leq 2 .$$

We give an example of a graph, its normalized Laplacian, and associated eigenvalues in Fig. 1. As we will later see, the eigenvalues λ_2 , ρ_2 , and μ_2 play a critical role in controlling expansion properties and defining Ramanujan graphs. In particular, the eigenvalue ρ_2 is called the *algebraic connectivity* of a graph. Due to its prevalence in the literature (see, for instance, [13,16,30]), we will choose to present our results in terms of this spectrum, keeping in mind that if G is k -regular, then

$$\rho_2 = k \cdot \mu_2 = k - \lambda_2.$$

Before proceeding, we describe the spectra of two graphs: the path and the cycle graph. We highlight these graphs as they are frequently elemental to the design of fundamental topologies (e.g. the torus, mesh, and hypercube are all obtained via graph products of cycles or paths). Unsurprisingly, their spectra are highly structured.

- The path of length $n - 1$, denoted P_n , has $n - 1$ edges and n vertices and adjacency spectrum

$$2 \cos \left(\frac{\pi j}{n+1} \right) \text{ for } j \in \{1, \dots, n\}.$$

- If the path of length $n - 1$ is modified to add self-loops at each of the endpoints, denoted P'_n , the adjacency spectrum becomes

$$2 \cos \left(\frac{\pi j}{n} \right) \text{ for } j \in \{0, \dots, n-1\}.$$

- The cycle of length n , denoted C_n , has n edges and vertices and adjacency spectrum

$$2 \cos \left(\frac{2\pi j}{n} \right) \text{ for } j \in \{0, \dots, n-1\}.$$

Finally, we use standard asymptotic notation: a function $f(n) = O(g(n))$ if for all sufficiently large values of n there exists a positive constant c such that $|f(n)| \leq c \cdot |g(n)|$; similarly, we write $f(n) = \Omega(g(n))$ if $g(n) = O(f(n))$, and $f(n) = \Theta(g(n))$ if both $f(n) = O(g(n))$ and $f(n) = \Omega(g(n))$. Lastly, $f(n) = o(g(n))$ if $\lim_{n \rightarrow \infty} \frac{f(n)}{g(n)} = 0$.

2.1 Network properties

Graph eigenvalues are deeply related to a number of fundamental network properties. In the case of supercomputing topologies, two such properties linked to communications performance are graph *diameter* and *bisection bandwidth*. Diameter (the maximum distance between vertices) is critical for latency, while bisection bandwidth (the minimum number of edges crossing a balanced bipartition of the vertices) measures the networks “bottleneckedness,” impacting all-to-all communication performance.

A plethora of work has shown both of these core network properties to be bounded and thus controlled by graph eigenvalues [20]. In particular, the eigenvalues of interest are the *spectral gap*, the difference in the largest two adjacency eigenvalues, or the *algebraic connectivity*, the second smallest Laplacian eigenvalue. For example, Alon and Milman [8] showed that the diameter is at most roughly $C \cdot \log n$, where C depends on algebraic connectivity and the maximum degree. More precisely,

Theorem 1 [8] *Let G be an n -vertex graph with algebraic connectivity ρ_2 and maximum degree Δ . Then*

$$\text{diam}(G) \leq 2 \left\lceil \sqrt{\frac{2\Delta}{\rho_2}} \log_2 n \right\rceil$$

A lower bound on graph diameter in terms of algebraic connectivity may also be obtained. For example, McKay [48] showed $\text{diam}(G) \geq \frac{4}{n\rho_2}$. In addition to these bounds on the maximum distance between vertices, *average* distance is upper and lower bounded in terms of algebraic connectivity as well; see [48]. Next, algebraic connectivity provides guarantees on minimum bisection bandwidth, as shown by Fiedler [24].

Theorem 2 [24] *Let G be an n -vertex graph with algebraic connectivity ρ_2 and bisection bandwidth $\text{BW}(G)$. Then*

$$\text{BW}(G) \geq \frac{\rho_2 n}{4}.$$

By considering Cheeger’s inequality [39,59], one can also obtain upper bounds on the bisection bandwidth in terms of ρ_2 for regular graphs.

Theorem 3 *For a connected k -regular, n -vertex graph G with algebraic connectivity ρ_2 , the bisection bandwidth satisfies*

$$\text{BW}(G) \leq \frac{\sqrt{2k\rho_2} \cdot kn}{2}.$$

We note that when ρ_2 is large this upper bound is quite loose. In fact, if G has m edges, an easy application of the first moment method [9] shows the bisection bandwidth is at most $\frac{m}{2}$. Note that if G is k -regular and ρ_2 is asymptotically k , then this first moment calculation shows that Theorem 2 is essentially tight and the bisection bandwidth is $\frac{kn}{4} (1 + o(1))$. Consequently, it can be shown that Ramanujan graphs (defined in Sect. 3) have nearly optimal bisection bandwidth among all k -regular graphs.

Lastly, we note that algebraic connectivity provides bounds on edge and vertex *connectivity*, the minimum number of edges and vertices that must be deleted in order to disconnect the graph, respectively. In the context of computer interconnection networks, vertex connectivity is often referred to as *fault tolerance* (e.g., [5]); more precisely, fault tolerance is defined as one less than vertex connectivity. Denoting

vertex and edge connectivity as $\kappa(G)$, $\kappa'(G)$, respectively, it is obvious that $\kappa(G) \leq \kappa'(G) \leq \Delta(G)$. Fiedler [24] proved

$$\kappa(G) \geq \rho_2,$$

and hence, larger algebraic connectivity guarantees more robust fault tolerance. For more spectral bounds on vertex and edge connectivity, the reader is referred to [4] and for further a more complete survey of the relationship between algebraic connectivity and numerous graph invariants, see [49]. Such spectral bounds have practical utility: for a number of graph topologies, exact diameter, bisection bandwidth, etc. may be unknown or difficult to compute, and hence, eigenvalues may serve as a proxy. In summary, the bounds we've reviewed motivate algebraic connectivity as a key parameter of interest intimately related to a plethora of structural properties important to interconnection network design. In the next section, we define graphs with optimal spectral gap, known as *Ramanujan graphs*, and discuss their expansion properties.

3 Ramanujan graphs

Ramanujan graphs are regular graphs with nearly optimal expansion properties. Loosely speaking, expansion means that every “not too large” set of vertices has a “not too small” set of neighbors. One way of measuring such expansion is the *vertex isoperimetric number* of a graph, given by

$$h(G) = \min_{\substack{X \subseteq V(G) \\ 2|X| \leq |V(G)|}} \frac{|\partial X|}{|X|},$$

where ∂X denotes the neighbors of vertices in X that are not in X . We illustrate examples of vertex isoperimetric ratios of sets in Fig. 2. This notion of expansion, as well as others such as the *edge isoperimetric constant*, has been shown to be intimately related to the second largest adjacency eigenvalue of a graph. For example, Tanner [61] proved a lower bound on $h(G)$ in terms of this eigenvalue λ_2 , for a k -regular graph; namely,

$$h(G) \geq 1 - \frac{k}{2k - 2\lambda_2}.$$

Conversely, Alon and Milman [8] proved an upper bound on λ_2 in terms of $h(G)$:

$$k - \lambda_2 \geq \frac{h(G)^2}{4 + 2 \cdot h(G)^2}.$$

Putting these two bounds together, it is clear that smaller values of λ_2 yield larger values of $h(G)$ and hence better expansion. Other bounds, such as Cheeger's inequality and Buser's inequality [18], similarly tie eigenvalues to other notions of expansion, like the Cheeger constant. Given the breadth of expansion properties reflected through




Fig. 2 Examples of a vertex set X , in red, with vertex boundary ∂X , in green. For **a**, the set has vertex isoperimetric ratio 4. For **b**, this ratio is $2/3$, which is the minimum over all subsets whose size doesn't exceed half the vertices, and thus is the graph's vertex isoperimetric number (color figure online)

eigenvalues, it is natural to measure expansion directly in terms of the spectra itself. Accordingly, researchers have sought *spectral expanders*, families of graphs with small λ_2 . The most well-known such family is called *Ramanujan graphs*.

Definition 1 A k -regular graph G is called Ramanujan if

$$\lambda(G) \leq 2\sqrt{k - 1},$$

where $\lambda(G)$ denotes the largest magnitude adjacency eigenvalue of G not equal to $\pm k$.

Ramanujan graphs are, in a sense, optimal spectral expanders since they achieve the asymptotic theoretical minimum given by Alon–Boppana theorems. The Alon–Boppana theorem [6,51] states that for a k -regular graph with second largest (in magnitude) adjacency eigenvalue λ and diameter D , we have

$$\lambda \geq 2\sqrt{k - 1} \left(1 - \frac{2}{D} \right) - \frac{2}{D}.$$

As an immediate corollary, if $(G_i)_{i=1}^\infty$ is a family of connected, k -regular, n -vertex graphs with $n \rightarrow \infty$ as $i \rightarrow \infty$, then

$$\liminf_{i \rightarrow \infty} \lambda(G_i) \geq 2\sqrt{k - 1}.$$

Hence, we see that Ramanujan graphs attain the theoretical asymptotic optimum spectral expansion. While the Alon–Boppana theorem pertains to regular graphs, variants of the theorem have been proposed for the case of irregular graphs, see [19,33,67].

As a consequence of their optimal spectral expansion, Ramanujan graphs possess beneficial structural properties via the bounds mentioned in Sect. 2. In particular, not only does the Ramanujan property guarantee at least nearly optimal bisection bandwidth, but also controls the number of edges between *any* collection of vertices, not just

bisections. This stronger property is known as the discrepancy property [21]. Specifically, using tools of spectral graph theory, if G is an n -vertex k -regular Ramanujan graph, we have that for any two sets of vertices X and Y ,

$$\left| e(X, Y) - \frac{k}{n} |X| |Y| \right| \leq \frac{2\sqrt{k-1}}{n} \sqrt{|X|(n-|X|)|Y|(n-|Y|)},$$

where $e(X, Y)$ is the number of edges between the sets X and Y . Roughly speaking, this says that in any Ramanujan topology the number of edges between two sets scales roughly like the expected number of edges between two sets in a similarly dense random graph. In particular, if a process is active on αn fraction of the nodes of the supercomputing topology, then bisection bandwidth on the active nodes is at least

$$\frac{\alpha kn}{2} \left(\frac{\alpha}{2} - \frac{2\sqrt{k-1}}{k} \left(1 - \frac{\alpha}{2} \right) \right)$$

independently of which αn nodes are chosen.

3.1 Ramanujan constructions

Providing explicit constructions of Ramanujan graphs is challenging. The first explicit constructions of Ramanujan graphs were given by Lubotzky, Phillips, and Sarnak [42], as well as independently by Margulis [45]. Both constructions are Cayley graphs that rely heavily on number-theoretic methods; indeed, the name ‘‘Ramanujan graph’’ was derived due to the application of the Ramanujan–Petersson conjecture from number theory in the aforementioned construction [42]. Below, we briefly describe and compare some of these constructions.

3.1.1 Lubotzky, Phillips, Sarnak construction

Definition 2 (*LPS Graphs*) The LPS graph $X^{p,q}$ is a $(q+1)$ -regular Cayley graph, defined for distinct primes p and q such that $p, q \equiv 1 \pmod{4}$. Letting i be any integer such that $i^2 \equiv -1 \pmod{p}$, the generating set S of $X^{p,q}$ is given by

$$S = \left\{ \left[\begin{array}{cc} \alpha_0 + i\alpha_1 & \alpha_2 + i\alpha_3 \\ -\alpha_2 + i\alpha_3 & \alpha_0 - i\alpha_1 \end{array} \right] \mid (\alpha_0, \alpha_1, \alpha_2, \alpha_3) \text{ is a solution of } \alpha_0^2 + \alpha_1^2 + \alpha_2^2 + \alpha_3^2 = q, \right. \\ \left. \alpha_0 > 1 \text{ is odd, and } \alpha_1, \alpha_2, \alpha_3 \text{ are even.} \right\},$$

and the group G of $X^{p,q}$ is

$$G = \begin{cases} \text{PSL}(2, \mathbb{F}_p) & \text{if } \left(\frac{q}{p}\right) = 1 \\ \text{PGL}(2, \mathbb{F}_p) & \text{if } \left(\frac{q}{p}\right) = -1 \end{cases},$$

where $\left(\frac{q}{p}\right)$ is the Legendre symbol.

We note that in the former case, the Cayley graph of $\text{PSL}(2, \mathbb{F}_p)$ with generating set S has $\frac{p(p^2-1)}{2}$ vertices and is non-bipartite, while in the latter case, the Cayley graph of $\text{PGL}(2, \mathbb{F}_p)$ with generating set S is bipartite with $p(p^2 - 1)$ vertices.

Using advanced number-theoretic techniques, Lubotzky, Phillips, Sarnak showed their construction has largest nontrivial adjacency eigenvalue at most $2\sqrt{q}$ and hence is Ramanujan. Additionally, they also showed their construction has other extremal combinatorial properties, such as having girth (i.e., the length of the shortest cycle) of $\Omega(\log_q n)$. From a computational standpoint, the LPS construction allows for explicit querying of vertex neighborhoods, which is a desirable property for analyzing exponentially large graphs.

The LPS construction may be used to generate infinite families of $(q + 1)$ -regular Ramanujan graphs, however, only for q prime with $q \equiv 1 \pmod{4}$ and n as function of p as given above. That is, despite having outstanding properties, the LPS construction is limited to Ramanujan graphs only of a certain degree—and for each such particular degree, only to a certain number of vertices n . In 1994, Morgenstern [50] partially ameliorated this restriction by extending the LPS construction to accommodate any prime power q , while showing this extended construction is still Ramanujan and satisfies all other combinatorial properties of the LPS graphs. Nonetheless, this still left open the general case of a given degree k and size n .

3.1.2 Marcus, Spielman, Snivrastra construction

In 2013 and 2015, Marcus, Spielman, Snivrastra gave new constructions of Ramanujan graphs using a new technique called the method of interlacing polynomials. Unlike the LPS construction, Marcus, Spielman, and Snivrastra's first construction [43] is valid for any given degree k , and second construction [44] is valid both for any k and number of vertices n . In both cases, their constructions can only be used to generate *bipartite* Ramanujan graphs.

While their interlacing family method implicitly suggests an algorithm to find an MSS graph, such an algorithm would require computing partially specified expected characteristic polynomials, for which no known polynomial time algorithms are known [23]. However, in [23], Cohen provided a polynomial time algorithm for computing such polynomials, thereby giving a deterministic algorithm that, for given degree k and even positive integer n , returns a bipartite Ramanujan graph, according to the construction given in [44], in polynomial time.

3.2 Related work in high-performance computing

Due to the aforementioned relationships between graph expansion and other properties important in network design, many proposed HPC network topologies consider graph expansion *implicitly*, making a comprehensive review of related work difficult. Before proceeding, we briefly survey related work that *explicitly* considers Ramanujan graphs or related expanders as network topologies in contexts pertinent to supercomputing. Perhaps most notably, in the context of datacenter architecture design, Valadarsky et al. [63] propose “Xpander,” which utilizes LPS graphs and the theory of graph lifts [14].

They evaluate Xpander theoretically, via simulation, and using a network emulator, finding that Xpander outperformed traditional data center designs; see [1] for more. In the early 1990s, Upfal [62] applied the theory of (α, β, n, d) -expander graphs to construct so-called multibutterfly networks. Later, Brewer, Chong, and Leighton [15] proposed a hierarchical expander construction, as a means to mitigate wiring complexity; they analyzed the fault tolerance of their so-called metabutterfly topology through simulation against the aforementioned multibutterfly. In optical network design, Paturi et al. [52] proposed using expander graphs to interconnect processors and subsequently analyze parallel algorithms for sorting, routing, associative memory, and fault tolerance. Lastly, in the context of sensor networks, Kar and Moura [36] propose using Ramanujan LPS graphs as communication networks supporting distributed decision making and test their performance on the convergence speed of distributed consensus.

4 Spectral gap in supercomputing topologies

Here, we survey a variety of supercomputing topologies. In addition to giving formal, and in some cases new or generalized, descriptions of the underlying graphs, we focus on analyzing their spectral gap, bisection bandwidth, and diameter. We first consider grid-like and grid variant topologies: the hypercube, generalized grid, torus, butterfly, Cube-Connected Cycles, and Data Vortex. Then, we consider several miscellaneous topologies: the CLEX, DragonFly, G -connected- H , and SlimFly topologies. Our results on algebraic connectivity and bisection bandwidth are summarized in Table 1.

Before proceeding, we first establish a key algebraic tool that we utilize frequently, allowing us to compute subsets of a given graphs spectra through that of a simpler, “reduced” graph.

Lemma 1 (Reduction Lemma) *Let G be a graph and let Γ be a subgroup of $\text{Aut}(G)$, the automorphism group of G . Let H be a weighted, directed, looped graph with vertex set given by the orbits of Γ over G and where the weight of edge from orbit σ to orbit τ is the total weight of an arbitrary vertex v in the orbit σ to the orbit τ . The spectrum of H is a subset of the spectrum of G . Furthermore, any eigenpair (λ, v) of G such that λ is not an eigenvalue of H has the property that v sums to zero along orbits of Γ .*

Proof Let (λ, w) be a right eigenpair of H . We define the vector w^Γ as follows: for any vertex v in G , define $e_v^T w^\Gamma = e_\sigma^T w$ where σ is the orbit containing v . Now let S be the collection of orbits and suppose v is in orbit τ . We then have that

$$\begin{aligned} e_v^T A_G w^\Gamma &= \sum_j e_v^T A_G e_j e_j^T w^\Gamma \\ &= \sum_{\sigma \in S} \sum_{j \in \sigma} e_v^T A_G e_j e_j^T w^\Gamma \\ &= \sum_{\sigma \in S} \sum_{j \in \sigma} e_v^T A_G e_j e_\sigma^T w \end{aligned}$$

Table 1 Table of bounds on algebraic connectivity and bisection bandwidth for common supercomputing topologies

	Nodes	Radix	ρ_2 upper bound	BW upper bound	Ramanujan ρ_2	Ramanujan BW
Butterfly(s, k)	sk^s	$2k$	$\frac{4\pi^2 k}{n^2}$	$\frac{(k+1)k^s}{2}$	$2k - 2\sqrt{2k - 1}$	$\frac{sk^{s+1}}{2} - \mathcal{O}(sk^{s+1/2})$
CCC(d)	$d2^d$	3	$\frac{\pi^2}{(d+2)^2}$	2^{d-1}	$3 - 2\sqrt{2}$	$\frac{3-2\sqrt{2}}{4}d2^d (1 + o(1))$
CLEX(k, ℓ)	k^ℓ	$2\ell k - k - 1$	$4k - 2$	$k^{\ell+1}$	$(1 - o(1))(2\ell k - 2\sqrt{2\ell k})$	$\frac{\ell k^{\ell+1}}{2} - \mathcal{O}(k^\ell \sqrt{k\ell})$
DataVortex(A)	$AC2^{C-1}$	4	$A2^{C-1}$	$A2^{C-2}$	$4 - 2\sqrt{3}$	$(1 - \sqrt{3}/2)AC2^{C-1} (1 + o(1))$
DragonFly(H)	$ H ^2 + H $	$\frac{2\ H\ }{ H } + 1$	$1 + \frac{1}{ H }$	$\left(\frac{ H +1}{2}\right)^2 + \text{BW}(H)$	$\frac{2\ H\ }{ H } + 1 - 2\sqrt{\frac{2\ H\ }{ H }}$	$\frac{(H +1)(2\ H\ + H)}{4} - \mathcal{O}\left(\sqrt{ H ^3 \ H\ }\right)$
$G \leftrightarrow_k H$	$ G H $	$\frac{2\ H\ }{ H } + k$	$k - \frac{k G }{2\ G\ } \lambda_2(G)$	$\frac{k G H }{2\ G\ } \text{BW}(G) + \text{BW}(H)$	$\frac{2\ H\ }{ H } + k - 2\sqrt{\frac{2\ H\ }{ H }} + k - 1$	$\frac{\ G\ (2\ H\ +k\ H\)}{4} - \mathcal{O}\left(G \sqrt{\ H\ \ H\ }\right)$
Hypercube(d)	2^d	d	2	2^{d-1}	$d - 2\sqrt{d - 1}$	$d2^{d-2} - \mathcal{O}(\sqrt{d2^d})$
PT(a, b), $a \geq b$	$10ab$	4	$\frac{4-3\cos(\frac{4\pi}{a})-\cos(\frac{2\pi}{a})}{5}$	$6b + ab + 5$	$4 - 2\sqrt{3}$	$ab(10 - 5\sqrt{3})(1 + o(1))$
SlimFly(q)	$2q^2$	$\frac{3q-1}{2}$	q	$\frac{q^3+q}{2}$	$\frac{3q-1}{2} - \sqrt{6q-6}$	$\frac{3q^3}{4} - \mathcal{O}(q^{3/2})$
Torus(k, d)	k^d	$2d$	$2(1 - \cos(\frac{2\pi}{k}))$	$2k^{d-1}$	$2d - \sqrt{8d-4}$	$dk^{d-1} - \mathcal{O}(\sqrt{dk^d})$




Fig. 3 An application of the reduction lemma to the fat tree

$$\begin{aligned}
 &= \sum_{\sigma \in S} e_{\tau}^T A_H e_{\sigma} e_{\sigma}^T w \\
 &= \lambda e_{\tau}^T w \\
 &= \lambda e_v^T w^{\Gamma}.
 \end{aligned}$$

As v is arbitrary, we have that λ is also an eigenvalue of G .

Now suppose that (λ, v) is an eigenpair for G and consider $v^{\Gamma} = \sum_{\sigma \in \Gamma} v_{\sigma}$. Since each σ is an automorphism of G , v_{σ} is also an eigenvector with eigenvalue λ . Thus, v^{Γ} is either an eigenvector with eigenvalue λ or it is the zero vector. Since v^{Γ} is constant over orbits of Γ , we can form v_H^{Γ} as the vector of values over orbits. It is clear that $A_H v_H^{\Gamma} = \lambda v_H^{\Gamma}$ and so either λ is in the spectrum of H or v_H^{Γ} is zero and v sums to zero over orbits of Γ . \square

We illustrate an application of the Reduction Lemma to a fat tree topology in Fig. 3. We note that the Reduction Lemma is almost certainly not new. In fact, it can be viewed as a special case of several other results on describing the interlacing of spectra of a matrix with a quotient matrix, see, for instance, [17, Chapter 1] and [16, Chapter 2].

It is worth noting that several topologies we will consider have implementations which have minor irregularities in the node radices. These deviations from regularity have little effect on the true performance of the network, and so we will add self-loops as needed to eliminate irregularity and simplify the analysis. This will not change the nature of any of our results as the bisection bandwidth and diameter both are unaffected by arbitrary self-loops.

4.1 Product (grid-like) topologies

For a number of high-dimensional supercomputing topologies, their underlying graphs can be obtained via repeated *graph products*. Product graphs are highly structured and possess properties which can sometimes be tightly controlled by those of their factor graphs. Below, we briefly describe three such topologies: the hypercube, torus, and generalized grid. These graphs are obtained via a particular graph product called the

Cartesian product, denoted $G \square H$. The graph $G \square H$ is on vertex set $V(G) \times V(H)$ and is defined by the edge condition: (u, u') and (v, v') are adjacent if and only if either

- $u = v$ and $\{u', v'\} \in H$, or
- $u' = v'$ and $\{u, v\} \in G$.

We note that the adjacency matrix of $G \square H$ can be written succinctly in terms of those of G and H ,

$$A_{G \square H} = A_G \otimes I + I \otimes A_H,$$

where I denotes the identity matrix and \otimes denotes Kronecker product. Using the above characterization, it is easy to show that the adjacency (or Laplacian) eigenvalues of $A_{G \square H}$ consists of $\lambda(G)_i + \lambda(H)_j$ over all $1 \leq i \leq |V(G)|$ and $1 \leq j \leq |V(H)|$; hence, the d -fold Cartesian product eigenvalues consist of all possible d -sums of the factor graph eigenvalues. In particular, the algebraic connectivity of $G \square H$ is the minimum of the algebraic connectivity of G and the algebraic connectivity of H .

Definition 3 (*Hypercube, Q_d*) The d -dimensional hypercube, Q_d , is on $n = 2^d$ vertices, defined by the d -fold Cartesian product $P_2 \square \dots \square P_2$, where P_2 is the path with 1 edge.

It is well known that Q_d has algebraic connectivity of 2 and bisection bandwidth $2^{d-1} = n/2$. The hypercube is a special case of the generalized grid graph, defined below.

Definition 4 (*Generalized Grid, G_{k_1, \dots, k_d}*) The d -dimensional, generalized grid the d -fold Cartesian product $P_{k_1} \square \dots \square P_{k_d}$, where P_{k_i} is the path of length $k_i - 1$.

We note that taking $d = 2, k_1 = m$, and $k_2 = n$ yields what is sometimes simply referred to as a grid graph, or $m \times n$ lattice, while taking $k_1 = \dots = k_d = 2$ yields Q_d . Using the aforementioned fact relating the Cartesian product eigenvalues to those of the factor graphs, it is easy to see the algebraic connectivity of G_{k_1, \dots, k_d} is $2 - 2 \cos(\pi / \max\{k_1, \dots, k_d\})$. Finally, we define the discrete torus topology, which is given by the cartesian product of cycles.

Definition 5 (*Torus, C_k^d*) The discrete torus C_k^d is the d -fold graph box product of a k -cycle, i.e., $C_k \square \dots \square C_k$. This graph is regular on $n = k^d$ vertices and has degree $2d$.

It is not difficult to show the algebraic connectivity of the torus C_k^d is $2(1 - \cos(2\pi/k))$.

4.2 Grid variants

The collection of topologies we consider in this section are closely related to topologies formed from the product operation, but with minor twists or modifications. Oftentimes, these topologies start from some grid-like layout and permute the connections or add small substructures to achieve desired properties.

4.2.1 Butterfly

One of the more well-known grid variants is the Butterfly topology [40]. In its most simple form the Butterfly topology consists of a series of shuffling layers based on the binary representation of the node names. More concretely, there are $p \log_2(p)$ switches arranged in a $\log_2(p)$ -by- p array of p switches in one of $\log_2(p)$ ranks. For each rank, each of the p switches is connected to two switches in the previous rank and two switches in the next rank. The nodes in rank i and position j are connected to the switch j and switch m in rank $i - 1$, where m is formed by flipping the i^{th} bit in binary representation of j . It is also connected to switch j and m' in rank $i + 1$, where m' is formed by switching the $(i + 1)^{\text{st}}$ bit in j . The Butterfly topology has diameter $\log_2(p)$ and bisection width $p/2$.

Definition 6 (*Butterfly, Butterfly*(k, s)) The k -ary, s -fly Butterfly network where there are s -layers of switches, and each switch has k “forward” connections. More concretely, the switches can be indexed by elements of $[s] \times [k]^s$. The “forward” connections from $(i, (a_1, \dots, a_s))$ to $(i + 1, (a'_1, \dots, a'_s))$ are formed by keeping all but the i th component of a fixed, i.e. $a_j = a'_j$ if $j \neq i$. Depending on the application, the s layers can either be connected linearly (no connection from layer s to layer 1), or cyclically (connection from layer s to layer 1). For convenience, we will restrict ourselves to the cyclic arrangement.

It is straightforward to see that these networks have a diameter of s by considering two elements in the same layer, (i, \mathbf{a}) and (i, \mathbf{b}) , where no coordinate of \mathbf{a} and \mathbf{b} agrees.

Proposition 1 Let G be a k -ary, s -fly Butterfly network. The bisection bandwidth of G is at most $\frac{(k+1)k^s}{2}$, and the algebraic connectivity is at most $2k - 2k \cos\left(\frac{2\pi}{s}\right)$.

Proof To upper bound the bisection bandwidth, we first consider the case where k is even and define $X = \left[\frac{k}{2}\right] \times [k]^{s-1}$. The bipartition we consider is then $([s] \times X, [s] \times \overline{X})$. In order for $(s, x) \in [s] \times X$ and $(s, x') \in [s] \times \overline{X}$ to be adjacent, it must be the case that $\{s, s'\} = \{1, 2\}$ and x and x' differ only in the first coordinate. This gives that $e([s] \times X, [s] \times \overline{X}) = 2 \left(\frac{k}{2}\right)^2 k^{s-1} = \frac{k^{s+1}}{2}$.

When k is odd, we construct a slightly more complicated partition. In particular, for $0 \leq i \leq s - 1$ define $X_i = \left\{\frac{k+1}{2}\right\}^i \times \left[\frac{k-1}{2}\right] \times [k]^{s-1-i}$ and let $X = \bigcup_i X_i$. Note that

$$|X| = \sum_{i=0}^{s-1} |X_i| = \frac{k-1}{2} \sum_{i=0}^{s-1} k^{s-1-i} = \frac{k-1}{2} \frac{k^s - 1}{k-1} = \frac{k^s - 1}{2}.$$

In particular, $([s] \times X, [s] \times \overline{X})$ is a bipartition of the vertex set of the k -ary, s -fly Butterfly network. Now to evaluate the bisection bandwidth we wish to count pairs $(u, v) \in X_i \times \overline{X}$ such that u and v differ in precisely one component. If we fix some $x_i \in X_i$, then for all $1 \leq j \leq i + 1$, modifying the j th component to be in $\frac{k+3}{2}, \dots, k$ yields such a pair. We note that modifying any entry $j > i + 1$ will preserve membership in X_i . Thus, the only remaining case to consider is when index

$(i + 1)$ is modified to have value $\frac{k+1}{2}$. This takes us outside the set X if and only if $x_i \notin \left\{\frac{k+1}{2}\right\}^i \times \left[\frac{k-1}{2}\right] \times \left\{\frac{k+1}{2}\right\}^j \times \left[\frac{k-1}{2}\right] \times [k]^{s-2-i-j}$ for some $0 \leq j \leq s - 2 - i$. As there are $\frac{k-1}{2} \frac{k^{s-i-1} + 1}{2}$ such terms x_i , this gives that the total number of pairs $(u, v) \in X_i \times \bar{X}$ which differ by exactly one component is given by

$$\begin{aligned} & \left(\frac{k-1}{2}\right) k^{s-i-1} (i + 1) \left(\frac{k-1}{2}\right) + \left(\frac{k-1}{2}\right) \left(\frac{k^{s-i-1} + 1}{2}\right) \\ &= \frac{(k-1)((i+1)k^{s-i} - ik^{s-i-1} + 1)}{4}. \end{aligned}$$

Thus, the bisection bandwidth is at most

$$\begin{aligned} & \sum_{i=0}^{s-1} 2 \frac{(k-1)((i+1)k^{s-i} - ik^{s-i-1} + 1)}{4} \\ &= \frac{s(k-1)}{2} + \frac{k-1}{2} \sum_{i=0}^{s-1} (i+1)k^{s-i} - ik^{s-i-1} \\ &= \frac{s(k-1)}{2} + \frac{k-1}{2} \left(k^s + 2\frac{k^s - k}{k-1} - (s-1)\right) \\ &= \frac{k-1}{2} + \frac{k^{s+1} - k^s}{2} + k^s - k \\ &= \frac{k^{s+1} + k^s - k - 1}{2}. \end{aligned}$$

To upper bound the algebraic connectivity, we note that there is an automorphism group of the Butterfly topology in which the orbits are given by the layers. Thus, by applying the reduction lemma to this automorphism group we get an s -cycle with edge multiplicity k . The bound on the algebraic connectivity then follows immediately. \square

4.2.2 Data Vortex

The Data Vortex topology was designed as a “streaming” topology with the idea that all of the data is constantly in motion (i.e., it is never buffered) and the data swirls from processors in the outer ring of the vortex toward the processors on the inside of the topology [31,58]. This streaming methodology has allowed the Data Vortex topology to handle the transmission of high volumes of data without suffering from significant congestion-related performance degradation (see [27,28,34,66] for a more in-depth discussion of the performance benefits of the Data Vortex topology.). Formally, the topology is defined a series concentric cylinders with “angular” transitions between them. Within the cylinders, there is a switching topology reminiscent of the layers of the 2-ary Butterfly topology. More concretely, we have:

Definition 7 (*Data Vortex, Data Vortex(A, C)*) The Data Vortex topology with parameters A, C is a graph with vertex set $\mathbb{Z}_A \times \mathbb{Z}_C \times \mathbb{Z}_2^{C-1}$ and edge set given by:

1. for all $(a, c, h) \in \mathbb{Z}_A \times \mathbb{Z}_C \times \mathbb{Z}_2^{C-1}$ there is an edge to $(a + 1, c + 1, h)$,
2. for all $(a, c, h) \in \mathbb{Z}_A \times (\mathbb{Z}_C - \{0\}) \times \mathbb{Z}_2^{C-1}$ there is an edge to $(a + 1, c, h + e_c)$ where e_c denotes the unit vector for the c^{th} component of \mathbb{Z}_2^{C-1} , and
3. for all $(a, c, h) \in \mathbb{Z}_A \times \{0\} \times \mathbb{Z}_2^{C-1}$ there is an edge to $(a + 1, c, h) = (a + 1, 0, h)$.

Although the Data Vortex is designed as a streaming topology (and is in particular, indirect), we will consider it as a direct topology in which each node denotes a compute node.

Proposition 2 *The algebraic connectivity of the Data Vortex topology with A angles and height H is at most $\min \left\{ 2 - 2 \cos \left(\frac{\pi}{C} \right), 2 - 2 \cos \left(\frac{2\pi}{A} \right) \right\} = \mathcal{O} \left(\frac{1}{\max \{ A^2, (C-1)^2 \}} \right)$. Furthermore, the bisection bandwidth is at most $A2^{C-2}$.*

Proof We begin by first noting that the vertices in the outer and inner ring of the Data Vortex have degree 3, so we will consider the topology formed by adding a self-loop to each of these vertices. Alternatively, we could add an edge between corresponding vertices in the inner and outer ring by observing that in typical use cases these vertices are connected to a common system, forming the “input” and “output” ports of the system. However, this modification results in essentially the same asymptotic behavior, so we choose the self-loop modification as it requires no assumptions about how the Data Vortex interacts with the processing layer of the overall system.

We first consider the bisection bandwidth by separating the vertices based on height, specifically partitioning into vertices of height $1, \dots, 2^{C-2}$, and those of height $2^{C-2} + 1, \dots, 2^{C-1}$. Clearly, this is a bisection. As no edge between concentric rings changes height, it suffices to consider only those edges internal to a ring. However, as only one ring flips the leading bit of the height vector, this gives that the bisection bandwidth is at most $A2^{C-2}$.

In order to bound the algebraic connectivity, we will apply the reduction lemma. Specifically, we consider the automorphism group generated by the bit-flip operations on the height. As these act uniformly on the height, the edges between successive rings are clearly preserved. Further, as the bit-wise differences are preserved by the bit-flip operations, this preserves edges on each ring. Under this automorphism group, the Data Vortex topology reduces to $C_A \square P'_C$ where P'_k is the k -vertex path with loops at each end. The result bounding the algebraic connectivity follows immediately. \square

4.2.3 Cube-Connected Cycles

Loosely speaking, the Cube-Connected Cycles (CCC) graph consists of a hypercube in which each vertex has been replaced by a cycle. Preparata and Vuillemin [54] proposed the Cube-Connected Cycles as a versatile network topology for connecting processors in a parallel computer, which emulates the robust connectivity properties of the hypercube, but (due to the cycle modification) only requires three connections per processor. They conclude that “by combining the principles of parallelism and pipelining, the CCC can emulate the cube-connected machine and shuffle-exchange network with no significant degradation in performance.”

As suggested in [56], CCC graph is a special case of a more general graph construction in which an arbitrary graph is connected in a hypercube structure. More precisely,

Definition 8 (*Cube-Connected Cycles, CCC(d)*) The d -dimensional cube-connected graph of a given graph G , denoted $CC(G, d)$, has vertex set $V(G) \times \{0, 1\}^d$ and edge condition $(v_i, x) \sim (v_j, y)$ if and only if

- $v_i \sim v_j$ in G , or
- $v_i = v_j$ and the hamming distance between x and y is 1.

Taking $G = C_d$ yields the well known *Cube-Connected Cycles* graph. Riess, Strehl, and Wanka proved the following result, which relates the characteristic polynomial of $CC(G)$ to those of loop-weighted variants of G :

Theorem 4 (Riess et al. [56]) *Let G be an d -vertex graph. For $s = (s_1, \dots, s_d) \in \{-1, 1\}^d$ and let $G[s]$ denote the graph obtained from G by adding a loop of weight s_i to each vertex i . Then,*

$$\chi(CC(G, d)) = \prod_{s \in \{-1, 1\}^d} \chi(G[s]),$$

where $\chi(G)$ denotes the characteristic polynomial of the adjacency matrix of G .

As an immediate consequence, we have that the spectral set of $CC(G, d)$ is the union of the spectral sets of $G[s]$ over all $s \in \{-1, 1\}^d$. Using their result, we can derive good estimates of the spectral expansion of the CCC. To do so, we first prove the following lemma.

Lemma 2 *Let G be a connected, n -vertex graph. The second largest adjacency eigenvalue of $CC(G, d)$ is the maximum eigenvalue of $G[s^*]$, where $s^* = (s_1, \dots, s_d) \in \{-1, 1\}^d$ is such that for some fixed $j \in [n]$, $s_j = -1$ and for all other $i \neq j$, $s_i = 1$.*

In the proof of Lemma 2, we will use the following basic fact:

Fact 1 *Let G be a connected, n -vertex graph. Let $r, t \in \{-1, 1\}^d$, $r \neq t$, be such that r agrees with t on any $i \in [n]$ where $t_i = 1$, and let $i_1, \dots, i_k \in [n]$ denote indices on which they differ, i.e., where $t_{i_j} = -1$ and $r_{i_j} = 1$ for $j \in [k]$. Then, the largest adjacency eigenvalue of $G[r]$ is strictly greater than that of $G[t]$.*

Proof Let A and A' denote the adjacency matrices of $G[r]$ and $G[t]$, respectively, and let $x = (x_1, \dots, x_n)$ denote the normalized, dominant eigenvector of $A' + I$, whose entries are all positive by the Perron–Frobenius theorem. By definition, we have $x^T(A + I)x - x^T(A' + I)x = 2 \sum_{j=1}^k x_{i_j}^2 > 0$. □

Proof (Proof of Claim 2)

From Theorem 4 and Fact 1, we have that the largest adjacency eigenvalue of $CC(G, d)$ is that of $G[\mathbf{1}_d]$ and furthermore that if $\mathbf{1}_d \neq t \in \{-1, 1\}^d$ does not satisfy the property in the claim, then there exists some s that does, which we denote s^* , such

that $\lambda_1(G[\mathbf{t}]) < \lambda_1(G[\mathbf{s}^*])$. So, let G and G' denote $G[\mathbf{1}_d]$ and $G[\mathbf{s}^*]$, respectively, on vertex set $\{v_1, \dots, v_n\}$, where v_j denotes the vertex in $G[\mathbf{s}^*]$ with a loop of weight -1 . Labeling the adjacency eigenvalues $\lambda_1 \geq \dots \geq \lambda_n$, all that remains to show is that

$$\lambda_2(G) < \lambda_1(G'). \quad (1)$$

By Cauchy's interlacing theorem, if we delete v_j from G and G' , we have

$$\begin{aligned} \lambda_2(G) &\leq \lambda_1(G \setminus v_j) \leq \lambda_1(G), \\ \lambda_2(G') &\leq \lambda_1(G' \setminus v_j) \leq \lambda_1(G'). \end{aligned}$$

But since $G \setminus v_j = G' \setminus v_j$, combining the above inequalities yields

$$\lambda_2(G) \leq \lambda_{\max}(G \setminus v_j) \leq \lambda_1(G').$$

To see the inequality in (1) is strict, assume for contradiction that $\lambda_1(G' \setminus v_j) = \lambda_1(G')$. Then, if $\mathbf{x} = (x_1, \dots, x_j, \dots, x_n)$ denotes the dominant eigenvector of A' associated with λ_1 , this implies if we set $x_j = 0$, the vector $(x_1, \dots, 0, \dots, x_n)$ is still an eigenvector of G' associated with λ_1 . But applying the Perron–Frobenius theorem to $A' + I$ yields that the dominant eigenvector of $A' + I$ (and hence that of A') is unique and has all entries positive, which is a contradiction. \square

Using Lemma 2, we have:

Proposition 3 *The algebraic connectivity of the d -dimensional, Cube-Connected Cycles is at most on the order of $2 \left(1 - \cos\left(\frac{\pi}{d+2}\right)\right)$.*

Proof By Lemma 2, it suffices to consider the largest adjacency eigenvalue λ_1 of the d -cycle with one loop of weight -1 on one vertex, and loops of weight 1 on other vertices. Letting A' denote this graphs adjacency matrix, a routine calculation shows that for $\mathbf{x} = (x_1, \dots, x_n)$ defined by $x_i = \sin\left(\frac{\pi i}{d+2}\right)$,

$$\lambda_1(A') \geq \frac{\langle \mathbf{x}, A'\mathbf{x} \rangle}{\langle \mathbf{x}, \mathbf{x} \rangle} = 2 \cos\left(\frac{\pi}{d+2}\right) + 1 + \frac{\sin^2\left(\frac{\pi}{d+2}\right) \left(2 \cos\left(\frac{\pi}{d+2}\right) - 2\right)}{\frac{d+1}{2} + \cos\left(\frac{2\pi}{d+2}\right)}.$$

We note that above expression is strictly larger than the second largest adjacency eigenvalue of the d -cycle with all loops of weight 1, for $d \geq 2$. \square

It is worth mentioning that the cube-connected graphs are really a specific instance of a more general technique of constructing supercomputing topologies, which we refer to as G -connected- H . As the generic G -connected- H topologies are not grid-like topologies, we will defer their discussion to Sect. 4.3.2.

4.3 Miscellaneous

In this section, we consider a few topologies that do not (necessarily) have a strong grid structure. Typically, these topologies have some sort of recursive or multi-layer structure in order to attempt to combine “good” properties of several types of graphs.

4.3.1 CLEX

“Clique-Expander” (CLEX) is a new supercomputing topology recently introduced by Lenzen and Wattenhofer [41]. The CLEX construction is recursive, starting with a specified number of cliques that are sparsely interconnected. According to the authors, the CLEX design is motivated by a desire to “localize the issue of an efficient communication network to much smaller systems which may reside on a single multi-core board.” CLEX is touted to have superior point-to-point communication properties, particularly when compared with toroidal topologies; nonetheless, the Lenzen and Wattenhofer acknowledge “the price we pay for these properties is [high] node degrees.”

In this section, we will define the CLEX topology and prove new bounds on the diameter, algebraic connectivity, and bisection bandwidth. As the authors of CLEX note that that “the high connectivity of a CLEX system could be considered an abstraction that can be replaced by any efficient local communication scheme within the cliques,” we generalize our spectral analysis of CLEX accordingly. In particular, our analysis allows one to replace the cliques of the CLEX construction with other graphs. We first begin by defining the CLEX graph, as given in [41].

Definition 9 (CLEX, $C(k, \ell)$) For given positive integers k and ℓ , a CLEX digraph, denoted $C(k, \ell)$, is on $n = k^\ell$ vertices with ℓ “levels” and is defined recursively. The base case is $C(k, 1) = K_k$, the complete graph on k vertices. The vertex set of $C(k, \ell + 1)$ is the $(\ell + 1)$ -fold cartesian product of $V(K_k)$. The edge set of $C(k, \ell + 1)$ consists of all edges from k copies of $C(k, \ell)$, with additional directed edges between these copies. Note the last entry in each vertex identifies which “copy” of $C(k, \ell)$ that vertex belongs to. The additional edges between these copies of $C(k, \ell)$ are given by the set:

$$\{((v_1, \dots, v_\ell, i), (v_1, \dots, v_{\ell-1}, j, v_\ell)) : i, j \in [k]\}.$$

With regard to the diameter of CLEX, the authors in [41] give an upper bound¹ of $C(k, \ell)$ as $2^\ell - 1$. We claim that the diameter is bounded by ℓ .

Proposition 4 *The diameter of the CLEX graph $C(k, \ell)$ is at most ℓ . Furthermore, this bound is tight.*

Proof We construct a walk of length ℓ between two arbitrary vertices of $C(k, \ell)$, (v_1, \dots, v_ℓ) and (w_1, \dots, w_ℓ) as follows:

¹ Note that there is actually a typo in their paper here, as they write $2^{1/\ell} - 1$, which is non-integer.

$$(v_1, \dots, v_\ell), (w_\ell, v_2, \dots, v_\ell), (w_1, w_\ell, v_3, \dots, v_\ell), \\ (w_1, w_2, w_\ell, v_4, \dots, v_\ell), \dots, (w_1, \dots, w_{\ell-1}, w_\ell).$$

Furthermore, this bound can be seen to be tight by considering the path between (i, i, \dots, i) and (j, j, \dots, j) for any $i \neq j$. Specifically, although each edge can modify up to two positions in the vector describing the vertex, it can change the count of any particular symbol in the string by at most one. \square

The rest of our analysis will consider the CLEX digraph as an undirected multi-graph (potentially with loops). Specifically, for every directed edge (i, j) in the CLEX digraph we will have an undirected edge $\{i, j\}$, and thus, the total degree of any vertex does not change. As our analysis only relies minimally on the structure of K_k , we will consider a generalized version of CLEX, denoted $C(G, \ell)$ where G is a t -regular, connected graph on k vertices. We note that both the regularity and connectivity conditions can be relaxed at various points in the following analysis; however, we make both assumptions for simplicity of presentation.

We first note that even when $G \neq K_k$, the arguments regarding the diameter follow exactly after accounting for the diameter of G and potentially directed nature of G .

Lemma 3 *Let G be a k -vertex graph, then*

$$C(G, \ell) = G \otimes I_{k^{\ell-1}} + \sum_{j=0}^{\ell-2} I_{k^j} \otimes M \otimes I_{k^{\ell-2-j}},$$

where $M \in \mathbb{Z}^{k^2 \times k^2}$ is given by

$$M_{(i,j),(a,b)} = \begin{cases} 2 & i = b, j = a \\ 1 & i = b, j \neq a \\ 1 & i \neq b, j = a \\ 0 & \text{otherwise.} \end{cases}$$

Proof The generic formula will follow immediately from the inductive characterization of the CLEX graphs. We note that the edges of $C(G, \ell + 1)$ can be partitioned in two sets: those that come from $C(G, \ell)$ and the cross-edges “between” copies of $C(G, \ell)$. Letting C_ℓ^G be the adjacency matrix for $C(G, \ell)$, the edge coming from the copies of $C(G, \ell)$ can be described by $C_\ell^G \otimes I_k$. Now note that an edge is added between $(v_1, \dots, v_{\ell-1}, v_\ell, v_{\ell+1})$ and $(w_1, \dots, w_{\ell-1}, w_\ell, w_{\ell+1})$ precisely when $v_i = w_i$ for $1 \leq i \leq \ell_i$ and $v_{\ell+1} = w_\ell$ or $w_{\ell+1} = v_\ell$. Thus, the cross-edges are given by $I_{k^{\ell-1}} \otimes M$ and we have that

$$C_{\ell+1}^G = C_\ell^G \otimes I_k + I_{k^{\ell-1}} \otimes M.$$

The non-inductive formula follows immediately from this relationship. \square

Lemma 4 Let $M \in \mathbb{Z}^{k^2 \times k^2}$ be defined by

$$M_{(i,j),(a,b)} = \begin{cases} 2 & i = b, j = a \\ 1 & i = b, j \neq a \\ 1 & i \neq b, j = a \\ 0 & \text{otherwise.} \end{cases}$$

We then have that $\text{spec}(M)$ is the multiset $\{2k, k^{(k-1)}, (-k)^{(k-1)}, 0^{(k-1)^2}\}$.

Proof Let $\{e_i\}$ be the standard basis vectors for \mathbb{R}^k and let $\mathbb{1}$ be the all ones vector in \mathbb{R}^k . We first note that

$$M = \sum_{i=1}^k (\mathbb{1} \otimes e_i) (e_i \otimes \mathbb{1})^T + (e_i \otimes \mathbb{1}) (\mathbb{1} \otimes e_i)^T.$$

It is easy to see at this point that $\mathbb{1} \otimes \mathbb{1}$ is an eigenvector of M with eigenvalue $2k$. Furthermore, we can see that the non-trivial eigenvectors must lie in $\text{span}\{\{\mathbb{1} \otimes e_i\}_i \cup \{e_i \otimes \mathbb{1}\}_i\}$, as a $2k - 1$ -dimensional subspace of \mathbb{R}^{k^2} .

Now consider

$$\begin{aligned} M(e_j \otimes \mathbb{1}) &= \sum_{i=1}^k \left((\mathbb{1} \otimes e_i) (e_i \otimes \mathbb{1})^T + (e_i \otimes \mathbb{1}) (\mathbb{1} \otimes e_i)^T \right) (e_j \otimes \mathbb{1}) \\ &= \sum_{i=1}^k (\mathbb{1} \otimes e_i) (e_i \otimes \mathbb{1})^T (e_j \otimes \mathbb{1}) + (e_i \otimes \mathbb{1}) (\mathbb{1} \otimes e_i)^T (e_j \otimes \mathbb{1}) \\ &= k(\mathbb{1} \otimes e_j) + \sum_{i=1}^k e_i \otimes \mathbb{1} \\ &= k(\mathbb{1} \otimes e_j) + \mathbb{1} \otimes \mathbb{1}. \end{aligned}$$

Similarly, we have that $M(\mathbb{1} \otimes e_j) = k(e_j \otimes \mathbb{1}) + \mathbb{1} \otimes \mathbb{1}$. From this, it is easy to see that

$$M(e_j \otimes \mathbb{1} - \mathbb{1} \otimes e_j) = -k(e_j \otimes \mathbb{1} - \mathbb{1} \otimes e_j),$$

for all j . Noting that $\sum_j (e_j \otimes \mathbb{1} - \mathbb{1} \otimes e_j) = 0$, we have that this yields a $k - 1$ -dimensional eigenspace associated with the eigenvalue $-k$. Finally, we note that $M(e_j \otimes \mathbb{1} + \mathbb{1} \otimes e_j - \frac{2}{k} \mathbb{1} \otimes \mathbb{1}) = k(e_j \otimes \mathbb{1} + \mathbb{1} \otimes e_j - \frac{2}{k} \mathbb{1} \otimes \mathbb{1})$; we similarly observe a $k - 1$ dimensional eigenspace associated with the eigenvalue k . As the dimension of the non-trivial eigenspaces is at most $2k - 1$, this provides a complete characterization of the spectrum. \square

Proposition 5 Let G be a t -regular, connected graph on k vertices. The algebraic connectivity of $C(G, \ell)$ is at most $t + 3k - 1$.

Proof First we note that since G is t -regular, $C(G, \ell)$ is $t + 2k(\ell - 1)$ regular. Now let (λ, v) be the eigenpair associated with the second largest eigenspace of G such that $\|v\| = 1$ and let $w = v \otimes \left(\frac{1}{\sqrt{k}} \mathbb{1}\right)^{\otimes \ell - 1}$. Since G is t -regular, we have that $\langle v, \mathbb{1} \rangle = 0$ and thus $\langle w, \mathbb{1}^{\otimes \ell} \rangle = 0$. Furthermore, since $\|v\| = 1$ and $\left\| \frac{1}{\sqrt{k}} \mathbb{1} \right\| = 1$, we have that $\|w\| = 1$. Thus, $w^T M w$ is a lower bound on the second largest eigenvalue of $C(G, \ell)$. We now note that

$$\begin{aligned} w^T M w &= w^T G \otimes I_{k^{\ell-1}} w + \sum_{i=0}^{\ell-2} w^T I_i \otimes M \otimes I_{\ell-2-i} w \\ &= \lambda + (v \otimes \mathbb{1})^T M (v \otimes M) + \sum_{i=1}^{\ell-2} 2k \\ &\geq \lambda - k + 2k(\ell - 2) \\ &= -1 - k + 2k(\ell - 2). \end{aligned}$$

Thus, the spectral gap is at most $t + 2k(\ell - 1) - (-1 - k + 2k(\ell - 2)) = t + 3k + 1$. □

Proposition 6 *Let G be a t -regular connected graph. If $\ell \geq 3$, the bisection bandwidth of $C(G, \ell)$ is at most $k^{\ell+1}$.*

Proof We may assume without loss of generality that the vertices of G are given by $[k] = \{1, \dots, k\}$, and thus, the vertex set of $C(G, \ell)$ is given by $[k]^\ell$. In order to upper bound the bisection bandwidth, we will provide two explicit partitions of the vertex set, one for the case when k is even and a modification construction for when k is odd. To that end, define A to be the set of odd integers in $[k]$ if k is even, and in $[k - 1]$ if k is odd. Similarly define A' to be the set of even integers in $[k]$. We note that if k is even, then $[k]$ is a disjoint union of A and A' , while if k is odd, $[k]$ is a disjoint union of A , A' , and $\{k\}$.

We first consider the case where k is even and define the sets $X = [k]^{\ell-2} \times (A \times A' \cup A' \times A)$ and $\bar{X} = [k]^{\ell-2} \times (A \times A \cup A' \times A')$. Since $|A| = |A'|$ and $[k]$ is a disjoint union of A and A' , it is clear that (X, \bar{X}) is a bisection of the $C(G, \ell)$.

Now let A_ℓ be the adjacency matrix of $C(G, \ell)$ and let $\mathbb{1}_X$ (respectively, $\mathbb{1}_{\bar{X}}$) be the indicator vector for the set X (respectively, \bar{X}). By definition

$$\text{BW}(C(G, \ell)) = \mathbb{1}_X^T A_\ell \mathbb{1}_{\bar{X}} = \mathbb{1}_X^T \left(G \otimes I_{k^{\ell-1}} + \sum_{j=0}^{\ell-2} I_{k^j} \otimes M \otimes I_{k^{\ell-2-j}} \right) \mathbb{1}_{\bar{X}}.$$

Noting that for any set S , we have that $\mathbb{1}_S^T I_{|S|} \mathbb{1}_{\bar{S}} = 0$ as S and \bar{S} are disjoint; this can be simplified to

$$\text{BW}(C(G, \ell)) = \mathbb{1}_X^T (I_{k^{\ell-3}} \otimes M \otimes I_k + I_{k^{\ell-2}} \otimes M) \mathbb{1}_{\bar{X}}$$

$$\begin{aligned}
 &= k^{\ell-3} \left(\mathbb{1}_{[k] \times A \times A'} + \mathbb{1}_{[k] \times A' \times A} \right)^T (M \otimes I_k + I_k \otimes M) \\
 &\quad \left(\mathbb{1}_{[k] \times A \times A} + \mathbb{1}_{[k] \times A' \times A'} \right) \\
 &= k^{\ell-3} \left(2 \frac{k}{2} \mathbb{1}_{[k] \times A}^T M \mathbb{1}_{[k] \times A'} + 4k \mathbb{1}_{A \times A'} M \mathbb{1}_{A \times A} \right),
 \end{aligned}$$

where the last line comes from the symmetry of A and A' and the symmetry of M in terms of the Kronecker product. Substituting in the definition for M , we get

$$\begin{aligned}
 (\mathbb{1} \otimes \mathbb{1}_A)^T M (\mathbb{1} \otimes \mathbb{1}_{A'}) &= \sum_{i=1}^k (\mathbb{1} \otimes \mathbb{1}_A)^T \left((e_i \otimes \mathbb{1})(\mathbb{1} \otimes e_i)^T + (\mathbb{1} \otimes e_i)(e_i \otimes \mathbb{1})^T \right) \\
 &\quad (\mathbb{1} \otimes \mathbb{1}_{A'}) \\
 &= \sum_{i=1}^k |A| k \left(e_i^T \mathbb{1}_{A'} \right) + k \left(e_i^T \mathbb{1}_A \right) |A'| \\
 &= 2k |A| |A'| \\
 &= \frac{k^3}{2}
 \end{aligned}$$

and

$$\begin{aligned}
 (\mathbb{1}_A \otimes \mathbb{1}_{A'})^T M (\mathbb{1}_A \otimes \mathbb{1}_A) &= \sum_{i=1}^k (\mathbb{1}_A \otimes \mathbb{1}_{A'}) \left((e_i \otimes \mathbb{1})(\mathbb{1} \otimes e_i)^T + (\mathbb{1} \otimes e_i)(e_i \otimes \mathbb{1})^T \right) \\
 &\quad (\mathbb{1}_A \otimes \mathbb{1}_A) \\
 &= \sum_{i=1}^k \left(\mathbb{1}_A^T e_i \right) \frac{k}{2} \frac{k}{2} \left(e_i^T \mathbb{1}_A \right) + \frac{k}{2} \left(\mathbb{1}_A^T e_i \right) \left(e_i^T \mathbb{1}_A \right) \frac{k}{2} \\
 &= \frac{k^3}{8}.
 \end{aligned}$$

Thus, we have that if k is even, the bisection bandwidth is $k^{\ell+1}$.

We now turn to the case where k is odd. Because of the parity issues in this case, it will be convenient to define the bipartition inductively. To that end, let (B, \bar{B}) be a bipartition of $C(G, \ell - 2)$ which witnesses the bandwidth such that $|B| + 1 = |\bar{B}|$. Now define the sets

$$\begin{aligned}
 Y &= [k]^{\ell-2} \times \left((A \times A') \cup (A' \times A) \cup (\{k\} \times [k - 1]) \right) \cup (B \times \{k\} \times \{k\}) \\
 \bar{Y} &= [k]^{\ell-2} \times \left((A \times A) \cup (A' \times A') \cup ([k - 1] \times \{k\}) \right) \cup (\bar{B} \times \{k\} \times \{k\}).
 \end{aligned}$$

It is clear that since $|A| = |A'|$ and $||B| - |\bar{B}|| = 1$, (Y, \bar{Y}) is a bipartition of $[k]^\ell$. Abusing notation slightly, we denote the set $[k]^{\ell-2} \times \left((A \times A') \cup (A' \times A) \right)$ by X and the set $[k]^{\ell-2} \times \left((A \times A) \cup (A' \times A') \right)$ as \bar{X} . If we again let A_ℓ denote the adjacency

matrix of $C(G, \ell)$, we have that

$$\mathbb{1}_X^T A_\ell \mathbb{1}_{\bar{X}} = (k-1)^3 k^{\ell-2}$$

by similar arguments as above. Additionally, we note that we have that

$$\begin{aligned} (\mathbb{1} \otimes \mathbb{1}_A)^T M (\mathbb{1} \otimes \mathbb{1}_{[k-1]}) &= k(k-1)^2 \\ (\mathbb{1} \otimes \mathbb{1}_A)^T M (\mathbb{1} \otimes e_k) &= k(k-1) \\ (\mathbb{1} \otimes \mathbb{1}_{[k-1]})^T M (\mathbb{1} \otimes e_k) &= 2k(k-1) \\ (\mathbb{1}_A \otimes \mathbb{1}_{A'})^T M (e_k \otimes \mathbb{1}_{[k-1]}) &= \frac{(k-1)^2}{4} \\ (\mathbb{1}_A \otimes \mathbb{1}_{A'})^T M (e_k \otimes e_k) &= 0 \\ (\mathbb{1}_{[k-1]} \otimes e_k)^T M (e_k \otimes e_k) &= k-1. \\ (\mathbb{1}_{[k-1]} \otimes e_k)^T M (e_k \otimes \mathbb{1}_{[k-1]}) &= k(k-1) \end{aligned}$$

Putting these calculations together, we get that the bandwidth of the partition (Y, \bar{Y}) is

$$(k-1)k^\ell + \mathbb{1}_{[k]^{\ell-2} \times [k-1]}^T (I_{k^{\ell-3}} \otimes M) \mathbb{1}_{B \times \{k\}} + \mathbb{1}_{B \times \{k\} \times \{k\}}^T A_\ell \mathbb{1}_{B \times \{k\} \times \{k\}}.$$

Observing that $A_\ell = A_{\ell-2} \otimes I_{k^2} + I_{k^{\ell-3}} \otimes M \otimes I_k + I_{k^{\ell-2}} \otimes M$, it is easy to see that

$$\mathbb{1}_{B \times \{k\} \times \{k\}}^T A_\ell \mathbb{1}_{\bar{B} \times \{k\} \times \{k\}} = \text{BW}(C(G, \ell-2)) + \mathbb{1}_{B \times \{k\}}^T (I_{k^{\ell-3}} \otimes M) \mathbb{1}_{\bar{B} \times \{k\}}$$

Now we note that terms involving $\mathbb{1}_{B \times \{k\}}$ sum to

$$\begin{aligned} &(\mathbb{1}_{[k]^{\ell-1}} - \mathbb{1}_{B \times \{k\}})^T (I_{k^{\ell-3}} \otimes M) \mathbb{1}_{B \times \{k\}} \\ &= \mathbb{1}_{[k]^{\ell-1}} (I_{k^{\ell-3}} \otimes M) \mathbb{1}_{B \times \{k\}} - \mathbb{1}_{B \times \{k\}}^T (I_{k^{\ell-3}} \otimes M) \mathbb{1}_{B \times \{k\}} \\ &= 2k|B| - \mathbb{1}_{B \times \{k\}}^T (I_{k^{\ell-3}} \otimes M) \mathbb{1}_{B \times \{k\}} \\ &\leq k(k^{\ell-2} - 1) \end{aligned}$$

Thus, we have that

$$\begin{aligned} \text{BW}(C(G, \ell)) &\leq (k-1)k^\ell + k^{\ell-1} - k + \text{BW}(C(G, \ell-2)) \\ &= k^{\ell+1} - k^\ell + k^{\ell-1} - k + \text{BW}(C(G, \ell-2)). \end{aligned}$$

Now, as $\text{BW}(C(G, 1)) \leq k^2$ and $\text{BW}(C(G, 2)) \leq k^3$, it is easy to see that by induction $\text{BW}(C(G, \ell)) \leq k^{\ell+1}$. \square

4.3.2 G -connected- H

The G -connected- H construction generalizes several different constructions, such as the Peterson Torus and Dragonfly topologies discussed in this section as well as the Cube-Connected Cycle topology discussed in Sect. 4.2.3. To see this, we first formally define what we mean by a G -connected- H topology.

Definition 10 (k -fold G -connected- H , $G \leftarrow_k H$) A k -fold G -connected- H topology, $\mathcal{G} = G \leftarrow_k H$, is constructed from a d -regular G and a r -regular td -vertex graph H . The vertex set of \mathcal{G} is $V_G \times V_H$ and $G[\{g\} \times V_H]$ is isomorphic to H for all vertices $g \in V_G$. The remaining edges form a k -regular graph on $V_G \times V_H$ satisfying that

$$e(\{v\} \times V_H, \{v'\} \times V_H) = \begin{cases} kt & \{v, v'\} \in E_G \\ 0 & \text{otherwise} \end{cases}.$$

When $k = 1$ will suppress the subscript and simply write $G \leftarrow H$.

Oftentimes, G is a Cayley graph and so the k -regular graph on $V_G \times V_H$ can be defined by a mapping from the generators of G to ordered pairs in V_H^2 . For example, we denote by Q_k the k -dimensional hypercube; we can view the Cube-Connected Cycle topology of Sect. 4.2.3 as a 1fold Q_k -connected- C_k . More concretely, we note that Q_k can be represented as the Cayley graph on \mathbb{Z}_2^k generated by the standard basis vectors, $\{e_1, \dots, e_k\}$. Since the generators of Q_k have order two, the matching edges can be formed by associating each generator with a fixed vertex of C_k .

This viewpoint can be extended to more complicated topologies, such as the Peterson Torus [35].

Definition 11 (*Peterson Torus, $PT(a, b)$*) Let $a, b \geq 2$ such that at least one of a or b is odd. Define the vertex set of the Peterson Torus Topology, $PT(a, b)$, as the set of ordered triples (x, y, p) where $0 \leq x < a$, $0 \leq y < b$, and $0 \leq p < 10$. Fixing the labels of the Peterson graph as given in Fig. 4a, the edge relationship is defined as:

- *internal edge* (x, y, p) is adjacent to (x, y, q) if p and q are adjacent in the Petersen graph.
- *longitudinal edge* $(x, y, 6)$ is adjacent to $(x, y + 1, 9)$.
- *latitudinal edge* $(x, y, 1)$ is adjacent to $(x + 1, y, 4)$.
- *diagonal edge* $(x, y, 2)$ is adjacent to $(x + 1, y + 1, 3)$.
- *reverse diagonal edge* $(x, y, 7)$ is adjacent to $(x - 1, y + 1, 8)$.
- *diameter edge* $(x, y, 0)$ is adjacent to $(x + \lfloor a/2 \rfloor, y + \lfloor b/2 \rfloor, 5)$.

This can be seen as a 1fold G -connected- H graph where G is the Cayley graph on $Z_a \times Z_b$ with generator set $\{\pm(0, 1), \pm(1, 0), \pm(1, 1), \pm(-1, 1), \pm(\lfloor a/2 \rfloor, \lfloor b/2 \rfloor)\}$ and H is the Peterson graph. We note that the condition that one of a or b is odd, is simply to ensure that the generator $(\lfloor a/2 \rfloor, \lfloor b/2 \rfloor)$ is not its own inverse, and so G has degree 10. By allowing multiple edges in G , this restriction can be eliminated.

We first consider the bisection bandwidth of $G \leftarrow_k H$. As this will depend explicitly on the number of nodes and edges in H , it is helpful to recall some standard notation

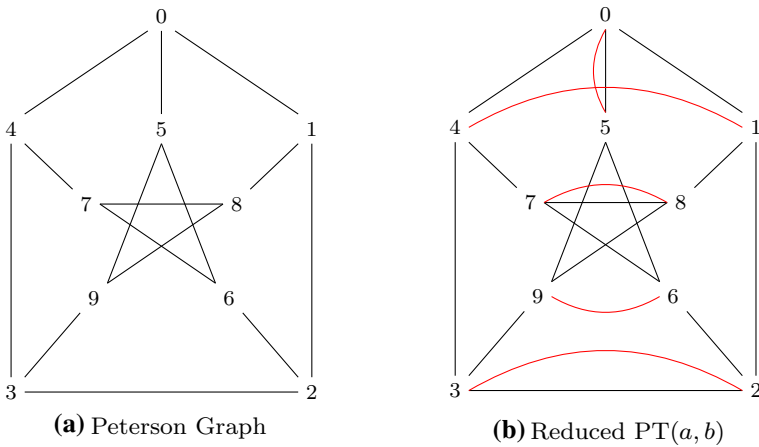


Fig. 4 Peterson graph and the Peterson graph with the symmetric function induced by the Peterson torus

first. Following the notation of West [65], the number of nodes in a network G will be denoted by $|G|$ and the number of edges will be denoted $\|G\|$.

Proposition 7 *Let $\mathcal{G} = G \leftarrow_k H$, then the bisection bandwidth of \mathcal{G} is at most $\frac{|G||H|}{2\|G\|}k \text{BW}(G) + \text{BW}(H)$.*

Proof We note that if $|G|$ is even, then the bipartition of G yielding $\text{BW}(G)$ lifts naturally to a bipartition of $G \leftarrow_k H$. As each edge in G is represented by $\frac{|G||H|}{2\|G\|}k$ edges in $G \leftarrow_k H$, this gives an upper bound of $\frac{|G||H|}{2\|G\|}k \text{BW}(G)$. If instead, $|G|$ is odd, the natural lift of the minimal bipartition doesn't yield a bipartition of $G \leftarrow_k H$. However, this can be corrected by splitting one of the copies of H , yielding the extra $\text{BW}(H)$ term. \square

We now turn the algebraic connectivity of $G \leftarrow_k H$. Because of the general structure of the matching edges and the potentially unstructured nature of G and H , the reduction lemma cannot be applied in general to G -connected- H graphs. However, there is still a natural symmetry formed by the G -connected- H structure, specifically the identification of vertices by common G labels or common H labels. However, because of the lack of automorphism structure we must turn to eigenvalue interlacing results such as the following by Haemmers.

Lemma 5 [17, Corollary 1.8] *Let A be an $n \times n$ real-symmetric matrix with eigenvalues $\lambda_1 \geq \lambda_2 \geq \dots \geq \lambda_n$. Let $\alpha_1, \dots, \alpha_m$ be a partition of the integers $\{1, 2, \dots, n\}$ into m nonempty consecutive sets of integers, where $|\alpha_i| = n_i$. Let A_{ij} be the submatrix of A defined by the entries whose row is in α_i and column is in α_j . Define B as the $m \times m$ real symmetric matrix with*

$$b_{ij} = \frac{\mathbb{1}_{n_i}^T A_{ij} \mathbb{1}_{n_j}}{n_i}.$$

The eigenvalues of B interlace the eigenvalues of A , in particular, $\lambda_2(B) \leq \lambda_2$.

Proposition 8 *Let G be a connected d -regular graph and let H be a connected r -regular, td -vertex graph, and let $\mathcal{G} = G \leftarrow_k H$ be a k -fold G -connected- H graph. Let λ_2 be the second largest eigenvalue of G , then the algebraic connectivity of \mathcal{G} is at most $k - \frac{k\lambda_2}{d}$.*

Proof Let A be the adjacency matrix of $\mathcal{G} = G \leftarrow_k H$. We will proceed to show that $\lambda_2(\mathcal{G}) \geq \frac{1}{d}\lambda_2 + r$, and then, the desired result follows immediately from the $(r + k)$ -regularity of \mathcal{G} . To this end, we will apply Lemma 5 to the partition of the vertices given by $\{\{v\} \times V_H\}_{v \in V_G}$. Abusing notation, for any $v, v' \in V_G$ we will denote by $A_{vv'}$ the submatrix induced by the rows $\{v\} \times V_H$ and columns $\{v'\} \times V_H$. Noting that $|\{v\} \times V_H| = td$ for all $v \in V_G$, we have that

$$\mathbb{1}_{td}^T A_{vv'} \mathbb{1}_{td} = \begin{cases} rtd & v = v' \\ tk & \{v, v'\} \in E_G \\ 0 & \{v, v'\} \notin E_G \end{cases}$$

and thus $B = rI + \frac{k}{d}A_G$ where A_G is the adjacency matrix of the graph G . The interlacing of the eigenvalues of B and \mathcal{G} provides the result immediately. \square

The strong dependence on the spectrum of G is unsurprising as the G -connected- H graphs implicitly inherit the connectivity structure of G , while increasing the relative degrees in a way that doesn't improve the spectral behavior of G . In particular, we note that another way of deriving Lemma 8 is to apply the Raleigh–Ritz formulation of $\lambda_2(\mathcal{G})$ and use the vector $\frac{1}{\sqrt{td}} \mathbb{1} \otimes w_2$ where (λ_2, w_2) is the second largest eigenpair of A_G .

It is natural to consider the implications of Lemma 5 when partitioning on the H -coordinate instead of the G -coordinate. Unfortunately, because of the unstructured nature of the k -regular graph, relatively little can be said. However, if the graph G is Cayley graph and the matching edges are tied to the generator set, then the automorphisms of G (specifically, those that follow from vertex transitivity of Cayley graphs) imply that there is an automorphism of $G \leftarrow H$ such that the orbits are given by $V_G \times \{h\}$ for $h \in V_H$. Then, the Reduction Lemma yields that there is multi-graph whose spectrum is a subset of the spectrum of $G \leftarrow_k H$. Specifically, take the graph H plus a k -regular graph (allowing self-loops) coming from the structure of the k -regular graph in $G \leftarrow_k H$. As an example, the Peterson Torus can be reduced with the reduction lemma to the graph illustrated in Fig. 4 with the red edges corresponding to the matching edges. Computing the algebraic connectivity of the reduced graph yields that the ρ_2 for the Peterson torus is at most 2. While this is small, it can be reduced further by applying Proposition 8.

Corollary 1 *Let $a \geq b \geq 2$ such that at least one of a or b is odd. The algebraic connectivity of $PT(a, b)$ is at most $\frac{4 - 3 \cos(\frac{4\pi}{a}) - \cos(\frac{2\pi}{a})}{5}$ and the bisection bandwidth is at most $6b + ab + 5$.*

Proof By Proposition 8 to bound the algebraic connectivity, it suffices to find the second largest eigenvalue of the Cayley graph G on the group $\Gamma = \mathbb{Z}_a \times \mathbb{Z}_b$ generated

by

$$S = \{\pm(1, 0), \pm(0, 1), \pm(1, 1), \pm(1, -1), \pm(\lfloor a/2 \rfloor, \lfloor b/2 \rfloor)\}.$$

Let $\chi : \Gamma \rightarrow \mathbb{C}^{|\Gamma|}$ be the character table for Γ . We recall that the spectrum of G is explicitly given by the multiset

$$\left\{ \sum_{s \in S} \chi_g(s) \mid g \in \Gamma \right\},$$

see, for instance, [16]. As Γ is the product of two cyclic groups, it is straightforward to explicitly determine the character table and get that the spectrum is given by the multiset

$$\left\{ \sum_{(s,t) \in S} e^{\frac{2\pi i x}{a} s} e^{\frac{2\pi i y}{b} t} \mid (x, y) \in \mathbb{Z}_a \times \mathbb{Z}_b \right\}.$$

In particular, this gives that λ_2 for G is given by

$$2 \max_{(x,y) \neq (0,0)} \cos\left(\frac{2\pi x}{a}\right) + \cos\left(\frac{2\pi y}{b}\right) + 2 \cos\left(\frac{2\pi x}{a}\right) \cos\left(\frac{2\pi y}{b}\right) + \cos\left(\frac{2\pi \lfloor a/2 \rfloor x}{a} + \frac{2\pi \lfloor b/2 \rfloor y}{b}\right).$$

It is relatively straightforward to see that the maximum is achieved when $(x, y) = (2, 0)$, yielding that λ_2 is at least

$$2 + 6 \cos\left(\frac{4\pi}{a}\right) + 2 \cos\left(\frac{4\pi \lfloor \frac{a}{2} \rfloor}{a}\right) \geq 2 + 6 \cos\left(\frac{4\pi}{a}\right) + 2 \cos\left(\frac{2\pi}{a}\right).$$

The upper bound on the bisection bandwidth will follow from Lemma 7. Specifically, as the Peterson torus is a $G \leftarrow H$ with G being the Cayley graph on $\mathbb{Z}_a \times \mathbb{Z}_b$ with generators S and H the Peterson graph, the bisection bandwidth is upper bounded by $\text{BW}(G) + \text{BW}(H)$. As the girth of the Peterson graph is 5, any collection of 5 vertices induces at most 5 edges. Thus, there are at least 5 edges crossing the cut, and this lower bound is achieved exactly by taking any of the 5 cycles in the Peterson graph.

For the bisection bandwidth of the graph on $\mathbb{Z}_a \times \mathbb{Z}_b$, we will denote the vertices by $[a] \times [b]$. If a is even, then the set $T = [\frac{a}{2}] \times [b]$ induces a bipartition with $6b + ab$ edges crossing the cut, that is, the edges corresponding to elements $\{\frac{a}{2}\} \times [b]$ and the generators $\{(1, -1), (1, 0), (1, 1)\}$, the edges corresponding to the elements $\{1\} \times [b]$ and the generators $\{(-1, -1), (-1, 0), (-1, 1)\}$, and the edges corresponding to an arbitrary vertex of T and the generators $\{(\lfloor \frac{a}{2} \rfloor, \lfloor \frac{b}{2} \rfloor), (-\lfloor \frac{a}{2} \rfloor, -\lfloor \frac{b}{2} \rfloor)\}$. In the case that a is odd, we consider the set $[\lfloor \frac{a}{2} \rfloor] \times [b] \cup \{\lceil \frac{a}{2} \rceil\} \times [\lfloor \frac{b}{2} \rfloor]$ and in a similar manner get that there are $6b + 2 \lfloor \frac{ab}{2} \rfloor$ edges crossing the cut, completing the proof. \square

4.3.3 DragonFly

As we will see, the DragonFly topology will end up being a specific class of G -connected- H topologies and can be understood in terms of the results of Sect. 4.3.2; however, due to their recent importance in “readily” available supercomputing topologies [3,10] we address them separately in this section. The motivating idea behind the DragonFly topology is to maximize the performance of a supercomputing topology while minimizing the overall cost of the system. To that end, Kim, Dally, Scott, and Abts designed the DragonFly topology around a two-level hierarchy [37]. The top level network employs an optical network to communicate over long distances (i.e. across the physical layout of the supercomputer), while the second layer employs an electrical network to communicate short distances (i.e. intrarack communication) and reduce the overall cost. While the specifications of Kim, et al. allow for arbitrary topologies for both the optical and electrical portions of the topology, the typically implementation uses a fully connected optical network combined with some other network for the electrical network, oftentimes either fully connected or a Butterfly variant. For example, the Cray Slingshot interconnect (which is being used for NSERC’s Perlmutter system) uses 64 port switches to build a DragonFly topology based on all-to-all connections for both the optical and electrical networks.

Definition 12 (*DragonFly, DragonFly(H)*) If H is an n -vertex, r -regular graph, then the DragonFly topology with parameter H consists of $n + 1$ copies of H together with a matching such that each edge goes between distinct copies of H . Alternatively, $\text{DragonFly}(H)$ may be thought of as a 1-fold $K_n \leftarrow H$.

We note that since the DragonFly topology can be represented as $G \leftarrow H$ topology, we immediately have bounds on the algebraic connectivity and bisection bandwidth.

Corollary 2 Let H be a connected graph and let D be the DragonFly topology generated by H . The algebraic connectivity of D is at most $1 + \frac{|H|}{2\|H\|}$ and the bisection bandwidth is at most

$$\left(\frac{|H| + 1}{2}\right)^2 + \text{BW}(H).$$

Proof Noting that $D = K_{|H|+1} \leftarrow H$ and the second largest adjacency eigenvalue of the complete graph is -1 , the bound on the algebraic connectivity follows immediately from Proposition 8. To provide the upper bound on the bisection bandwidth, consider a equipartition of the $|H| + 1$ copies of H . If $|H|$ is odd, then the only edges crossing the partition are “matching” or “optical” edges and there are $\left(\frac{|H|+1}{2}\right)^2$ of them. However, if $|H|$ is even, then one of the copies of H must also be partitioned yielding

$$\left(\frac{|H|}{2}\right)^2 + \frac{|H|}{2} + \text{BW}(H) \leq \left(\frac{|H| + 1}{2}\right)^2 + \text{BW}(H)$$

edges crossing the partition. □

4.3.4 SlimFly

In [11], Besta and Hoefler suggested that it would be advantageous to consider topologies that have close to the maximum number of nodes for a given radix and diameter. The upper bound on the number of nodes of a k -regular graph of diameter d is given by $1 + k \sum_{i=0}^{d-1} (k-1)^i$ and is referred to as the Moore bound. The class of graphs exactly achieving this bound, known as Moore graphs, has been extensively studied and shown to have significant limitation on both the radix and size, see [47].

In this context, Besta and Hoefler propose the SlimFly topology based on the construction of McKay, Miller, and Širáň [46] which is close to achieving the Moore bound. These SlimFly topologies have a single parameter q , which is a prime power such that $q \equiv 1 \pmod{4}$ and results in a topology on $2q^2$ nodes with degree $\frac{3q-1}{2}$.

Definition 13 (*SlimFly*, *SlimFly*(q)) Let ζ be a primitive q th-root of unity over the Galois field \mathbb{F}_q . The vertices are then elements of $\{0, 1\} \times \mathbb{F}_q \times \mathbb{F}_q$. The edge set is broken into three sets:

1. $\{(0, x, y), (0, x, y')\}$ where $y - y' = \zeta^i$ and $i \equiv 0 \pmod{2}$,
2. $\{(1, m, c), (1, m, c')\}$ where $c - c' = \zeta^j$ and $j \equiv 1 \pmod{2}$, and
3. $\{(0, x, y), (1, m, c)\}$ where $y = mx + c$.

Proposition 9 Let q be a prime-power such that $q \equiv 1 \pmod{4}$. The algebraic connectivity of the SlimFly topology with parameter q is q .

Proof In order to bound the algebraic connectivity, we will use the Reduction Lemma. To that end, let ζ be a primitive root of the Galois field \mathbb{F}_q and define γ_ζ by $(0, x, y) \mapsto (0, x, y + \zeta)$ and $(1, m, c) \mapsto (1, m, c + \zeta)$. It is easy to see that this is an automorphism of the SlimFly topology and that the orbits of the group generated by this automorphism are given by $\{0\} \times \{x\} \times \mathbb{F}_q$ and $\{1\} \times \{m\} \times \mathbb{F}_q$ for $x, m \in \mathbb{F}_q$. As an arbitrary element $(0, x, y) \in \{0\} \times \mathbb{F}_q \times \mathbb{F}_q$ has precisely one neighbor in the orbit $\{1\} \times \{m\} \times \mathbb{F}_q$ for any $m \in \mathbb{F}_q$, namely $(1, m, y - mx)$, we have that the reduction graph H is a complete bipartite graph $K_{q,q}$ with $\frac{q-1}{2}$ self-loops at every vertex. As the algebraic connectivity of this graph is q , by the Reduction Lemma we have that the algebraic connectivity of the is at most q .

Now we will show that the algebraic connectivity is exactly q . To this end, recall that the eigenspace associated with any eigenvalue that is not present in the spectrum of the reduced graph has the property that the entries sum to zeros over all of the orbits. That is, if v is such an eigenvector and $\mathbb{1}_\sigma$ is the indicator function for the orbit σ , then $v^T \mathbb{1}_\sigma = 0$. Furthermore, since the orbits of the automorphism are Cayley graphs on \mathbb{F}_q , the eigenvectors can be expressed in terms of the characters of $(\mathbb{F}_q, +)$. Additionally, the eigenvalues associated with χ_f are given by the character sums over the generators. Specifically, the eigenvalue associated to the non-trivial character χ on the Cayley graph is generated by ζ^{2j} is $\mu = \sum_{j=1}^{q-1/2} \chi(\zeta^{2j})$, while the eigenvalue associated with χ on the Cayley graph is generated by ζ^{2j-1} is $\sum_{j=1}^{q-1/2} \chi(\zeta^{2j-1}) = -1 - \mu$. Thus, let $\{\chi_f\}_{f=1}^{q-1}$ be the set of non-trivial characters of $(\mathbb{F}_q, +)$. The eigenvector v can then

be expressed as

$$\sum_{i \in \{0,1\}} \sum_{x \in \mathbb{F}_q} \sum_{f \in [q-1]} \frac{a_{i,x,f}}{\sqrt{q}} e_i \otimes e_x \otimes \chi_f$$

where $\sum_{i,x,f} \alpha_{i,x,f}^2 = 1$

Now letting A be the adjacency matrix of the SlimFly topology, we consider the quadratic form $\bar{v}^T A v$ in three parts. The portion corresponding to edges induced by $\{0\} \times \mathbb{F}_q \times \mathbb{F}_q$, the portion corresponding to edges induced by $\{1\} \times \mathbb{F}_q \times \mathbb{F}_q$, and the portion corresponding to edges between these two sets. It is easy to see that the contribution of the edges internal to these two sets is given by $\sum_{f=1}^{q-1} \sum_{x \in \mathbb{F}_q} \overline{\alpha_{0,x,f}} \alpha_{0,f,x} \mu_f$ and $\sum_{f=1}^{q-1} \sum_{m \in \mathbb{F}_q} \overline{\alpha_{0,m,f}} \alpha_{1,m,f} (-1 - \mu_f)$, respectively. Recalling the edges between the two sets are governed by the relationship $y = mx + c$ for $(x, y), (m, c) \in \mathbb{F}_q^2$, we have that the contribution of those edges to the quadratic form is

$$\begin{aligned} &\sum_{f,g=1}^{q-1} \sum_{(x,y) \in \mathbb{F}_q^2} \sum_{(m,c) \in \mathbb{F}_q^2} \frac{\overline{\alpha_{0,x,f}} \alpha_{1,m,g}}{q} \overline{\chi_f(y)} \chi_g(c) \mathbb{1}_{y=mx+c} \\ &\quad + \frac{\alpha_{0,x,f} \overline{\alpha_{1,m,g}}}{q} \chi_f(y) \overline{\chi_g(c)} \mathbb{1}_{y=mx+c}. \end{aligned}$$

Now we note that the nonzero entries in the sum occur when $c = y - mx$. Furthermore, χ_g is a homomorphism into (\mathbb{C}, \times) so $\chi_g(y - mx) = \chi_g(y) \chi_g(-mx)$. By additionally recalling that $\{\chi_f\}_{f=1}^{q-1}$ is an orthogonal basis, this sum simplifies to

$$\sum_{f=1}^{q-1} \sum_{x \in \mathbb{F}_q} \sum_{y \in \mathbb{F}_q} \sum_{m \in \mathbb{F}_q} \overline{\alpha_{0,x,f}} \alpha_{1,m,f} \chi_f(-mx) + \overline{\alpha_{1,m,f}} \alpha_{0,x,f} \overline{\chi_f(-mx)}.$$

Thus, letting M be the diagonal matrix formed from $\{\mu_f\}_{f=1}^{q-1}$, we have that the norm of the quadratic form is bounded above by the largest eigenvalue of

$$\mathcal{M} = \begin{bmatrix} M & I \\ I & -M - I \end{bmatrix}.$$

Motivated by this formulation, we consider the auxiliary problem

$$\max_{x^2+y^2=1} \mu x^2 + 2xy - (\mu + 1)y^2.$$

Noting that we may assume that $x, y \geq 0$, this can be reparameterized as

$$\max_{\delta \in [-1,1]} \mu \left(\frac{1}{2} + \frac{1}{2} \delta \right) + 2 \sqrt{\frac{1}{2} + \frac{1}{2} \delta} \sqrt{\frac{1}{2} - \frac{1}{2} \delta} - (\mu + 1) \left(\frac{1}{2} - \frac{1}{2} \delta \right).$$

The derivative of the objective function is $\frac{2\mu+1}{2} - \frac{\delta}{\sqrt{1-\delta^2}}$ with roots $\pm \frac{2\mu+1}{\sqrt{(2\mu+1)^2+4}}$. Thus, the largest eigenvalue of \mathcal{M} is

$$\max_{f \in [q-1]} \max \left\{ \mu_f, -\mu_f - 1, -\frac{1}{2} + \frac{1}{2}\sqrt{(2\mu_f + 1)^2 + 4}, -\frac{1}{2} - \frac{1}{2} \frac{(2\mu_f + 1)^2 - 4}{\sqrt{(2\mu_f + 1)^2 + 4}} \right\}.$$

Using the fact that the Cayley graph generated by the odd powers of ζ and the Cayley graph generated by the even powers of ζ are isomorphic (via $x \mapsto \zeta x$), this reduces to $-\frac{1}{2} + \frac{1}{2}\sqrt{(2\mu + 1)^2 + 4}$ where μ is the second largest eigenvalue of the Cayley graph generated by the even powers of ζ . Using the fact that this Cayley graph is edge transitive and has diameter 2, we get that $\mu \leq \frac{q-1}{2} - \frac{1}{4}\frac{q-1}{2}$ (see [21, Section 7.3]). Combining these results, we have that the largest eigenvalue not represented in the reduced graph is at most

$$-\frac{1}{2} + \frac{1}{2}\sqrt{\left(\frac{3}{4}q + \frac{1}{4}\right)^2 + 4} < \frac{q-1}{2}$$

for $q \geq 5$. □

Proposition 10 *Let q be a prime-power such that $q \equiv 1 \pmod{4}$. The bisection bandwidth of the SlimFly topology with parameter q is at most $\frac{q(q^2+1)}{2}$ and at least $\frac{q^3}{2}$.*

Proof Let $X \subset \mathbb{F}_q$ such that $|X| = \frac{q-1}{2}$ and let \bar{X} be the complement of X . We consider the bipartition $\{0\} \times X \times \mathbb{F}_q \cup \{1\} \times \bar{X} \times \mathbb{F}_q$. We note that there are no edges between $\{0\} \times X \times \mathbb{F}_q$ and $\{0\} \times \bar{X} \times \mathbb{F}_q$, and similarly, there are no edges between $\{1\} \times X \times \mathbb{F}_q$ and $\{1\} \times \bar{X} \times \mathbb{F}_q$. Now, as $\{0\} \times \{x\} \times \mathbb{F}_q$ has exactly one edge to $\{0\} \times \{m\} \times \mathbb{F}_q$ for every $x, m \in \mathbb{F}_q$. Thus, the bisection bandwidth of the SlimFly topology is at most $q \left(\frac{q-1}{2}\right)^2 + q \left(\frac{q+1}{2}\right)^2 = \frac{q(q^2+1)}{2}$.

The lower bound follows from Lemma 9 and the lower bound on the bandwidth based on the algebraic connectivity. □

It is worth mentioning that the gap between the bisection bandwidth achieved by a $\frac{3q-1}{2}$ regular graph on $2q^2$ vertices and the bisection bandwidth of the SlimFly topology could be attributed to fact that the SlimFly topology is not a Moore graph. In fact, it is straightforward to construct a bisection of a Moore graph whose bisection bandwidth asymptotically matches the known lower bounds on the bisection bandwidth of a similar Ramanujan graph.

Proposition 11 *Let G be a Moore graph with regularity q and girth $2d + 1$. The bisection bandwidth of G is at most $\frac{q}{2} + \frac{q^2}{4} (q - 1)^{d-1}$ if q is even and $q + \frac{q^2-1}{4} (q - 1)^{d-1}$ if q is odd.*




Fig. 5 Proportional bisection bandwidth for supercomputing topologies by number of compute nodes

Proof Fix an arbitrary vertex v in G and let its neighbors be w_1, \dots, w_q . Since the girth of Moore graph is $2d + 1$, the diameter is d . For $i \in [q]$ define W_i as the set of vertices whose shortest path to v goes through w_i and define $S_i \subset W_i$ as the vertices are at distance precisely d from v . Note that since G is a Moore graph, for any vertex $s \in S_i$ all the neighbors of s must be in distinct sets S_j where $j \neq i$.

Suppose first that q is even and consider the bipartition $\left(\left(\bigcup_{i=1}^{\frac{q}{2}} W_i \right) \cup \{v\}, \bigcup_{i=\frac{q}{2}+1}^q W_i \right)$. Now clearly each edge in each of the W_i 's does not cross the bipartition, and so the only edges we need concern ourselves with are those adjacent to v and those adjacent to vertices of S_i . Now as each vertex in S_i is adjacent to a vertex in each of the S_j 's except S_i , this implies that there are $\frac{q}{2} + \frac{q}{2} \sum_{i=1}^{\frac{q}{2}} |S_i| = \frac{q}{2} \left(1 + \frac{q}{2} (q - 1)^{d-1} \right)$ edges crossing the bisection.

The construction for q odd is similar to the one for q even, except rather than placing all of $W_{\frac{q+1}{2}}$ on one side of the partition, the partitioning procedure is done on the trees rooted at the vertices of distance 2 from v in $W_{\frac{q+1}{2}}$. □

5 Conclusion

We provide in Table 1 a summary of the results on the bisection bandwidth and algebraic connectivity of the topologies considered in this work. Additionally, for comparison we provide bounds on the bisection bandwidth and algebraic connectivity for a similarly sized Ramanujan topology. We focus on bisection bandwidth in our comparison, although we remind the reader the spectral results summarized in Table 1 also provide bounds on a plethora of other salient interconnection network properties (such as diameter, average distance, and fault tolerance) via the theorems mentioned in Sect. 2. As closer inspection of the table makes clear, for each of these topologies there is a significant gap between the achieved value and the minimum guaranteed to be achievable in an equivalent Ramanujan topology. However, assessing these results *across* families is more challenging due to different input parameters and parameter multiplicities for each topology. To better enable such a comparison, in Fig. 5 we plot

the proportional² bisection bandwidth by number of compute nodes for each topology, as well as the minimum guaranteed by a Ramanujan topology. In general, the solid lines represent those topologies with switches comparable to current topologies (that is, having radix at most 64 as in the Cray Slingshot Topology [3,10] while the dashed lines represent the proportional bisection bandwidth achievable with next-generation switches (radix at most 128), and the dotted line represents those topologies that would require even higher radix switches. We note that even the limitations on the radix are not sufficient to uniquely determine the highest bisection bandwidth proportion for some topologies. Thus, we will also impose following additional assumptions on the topologies with an aim of avoiding trivial instantiations of the topology:

- *Butterfly* for the Butterfly topology, we assume that there are at least 3 ranks of switches, i.e. $s \geq 3$,
- *CLEX* for the CLEX topology we assume that there are at least two layers $\ell \geq 2$ and that the initial generating graph is the complete graph on at least 3 vertices,
- *Data Vortex* for the Data Vortex, we assume that there are at least 3 “cylinders,” i.e. $C \geq 3$,
- *DragonFly–Butterfly* similar to the Butterfly topology, for the DragonFly topology where the electrical network is given by a Butterfly network, we assume that $s \geq 3$, and
- *Torus* for the torus topology we assume that all the cycles are non-degenerate, i.e., that $k \geq 3$.

Even as we compare these upper bounds on the best possible bisection bandwidth for each topology against the worst possible in a Ramanujan topology, we still observe a sizable gap, with the 128 radix SlimFly and CLEX topologies the closest to the Ramanujan lower bound. We suspect the region where CLEX outperforms Ramanujan graphs is an artifact of the looseness of the analysis of CLEX for small parameter settings, rather than a true reflection of the relative sizes of the bisection bandwidths.

In light of the beneficial structural properties of random graphs, it is natural to ask whether any potential utility of Ramanujan supercomputing topologies is already offered by randomized constructions, such as the well-known Jellyfish topology. Indeed, such topologies are touted for their low diameter, short average path lengths, and high bisection bandwidth [60]. Although random regular graphs are not quite Ramanujan, it is true that random d -regular graphs have good spectral expansion. Notably, Friedman’s celebrated proof [25] of Alon’s second eigenvalue conjecture [6] showed that if G is a random k -regular graph on n vertices, then with probability going to 1 as $n \rightarrow \infty$, we have $\lambda(G) \leq 2\sqrt{k-1} + o(1)$. Thus, in the limiting sense, random regular graphs are “almost Ramanujan.” Nonetheless, randomized constructions are also limited as interconnection topologies in that they pose serious challenges for routing, physical layout, and wiring [60]. In these regards, structured topologies offer advantages.

Consequently, one may ask whether more structured families, such as Cayley graphs, might serve as a more amenable alternative to random constructions. Since many of the popular topologies can be phrased as Cayley graphs (e.g., the torus and

² Relative to sum of the graph degrees, or twice the number of links.

hypercube topologies) or have a strong connections to Cayley graphs (e.g., the Slim-Fly and Peterson torus topologies), it is natural to speculate that a Cayley graph could serve as the basis of a strong supercomputing topology. Indeed, work [5] investigating Cayley graphs as interconnection networks dates back to at least the 1980's, see [32] for a survey. In particular, *abelian* Cayley graphs may seem particularly promising because the classification of abelian groups gives a natural means of easily performing efficient routing. However, abelian Cayley graphs do not offer the spectral expansion of Ramanujan graphs: as a consequence of a result of Cioabă [22], there is a constant $C(k, \epsilon)$ such that if the group has more elements than $C(k, \epsilon)$, then any Cayley graph generated by a k -element set has algebraic connectivity at most ϵ . Thus, for any fixed radix k , there does not exist an infinite family of radix k abelian Cayley graphs which are Ramanujan.

Given these trade-offs between randomized designs and highly structured Cayley graph designs, we believe the explicit Ramanujan construction by Lubotsky, Phillips, and Sarnak warrants further investigation as a candidate for supercomputing interconnection networks. By virtue of their optimal spectral expansion, LPS graphs offer many of the same (if not better) structural properties exhibited by random regular graphs. Yet, as highly structured Cayley graphs, LPS graphs may be more amenable to practical considerations and easier to develop efficient routing schemes for than random constructions. Indeed, recent work by Sardari [57], as well as Pinto and Petit [53] investigating short paths in LPS graphs, shows that, while sometimes challenging to analyze, the local structure of these topologies may be exploited for the purposes of routing. While the work we've done here attests to the structural benefits of LPS graphs over other supercomputing topologies, additional work is needed to better assess the benefits of utilizing LPS graphs as interconnection networks in practice.

Acknowledgements We would like to thank Andres Marquez, Kevin Barker, and Carlos Ortiz-Marrero for helpful discussions. This work was supported by the High Performance Data Analytics program at PNNL. Information Release PNNL-SA-147472.

References

1. *Xpander project page*. <https://husant.github.io/Xpander/>
2. Bob Metcalfe eats his words (1997) Internet Computing Online 1
3. *Slingshot: The interconnect for the exascale era*. Technical Report, Cray Inc., February (2019)
4. Abiad A, Brimkov B, Martinez-Rivera X, Suil O, Zhang J (2018) Spectral bounds for the connectivity of regular graphs with given order. *Electron J Linear Algebra* 34:428–443
5. Akers SB, Krishnamurthy B (1987) On group graphs and their fault tolerance. *IEEE Trans Comput* C-36:885–888
6. Alon N (1984) Eigenvalues and expanders. *Combinatorica* 6 (1986) 83–96. *Theory of computing* (Singer Island, Fla., 1984)
7. Alon N, Chung FRK, Graham RL (1994) Routing permutations on graphs via matchings. *SIAM J Discrete Math* 7:513–530
8. Alon N, Milman V (1985) λ_1 , isoperimetric inequalities for graphs, and superconcentrators. *J Comb Theory Ser B* 38:73–88
9. Alon N, Spencer JH (2000) *The probabilistic method*. Wiley-Interscience series in discrete mathematics and optimization. Wiley-Interscience (John Wiley & Sons), New York, second ed. With an appendix on the life and work of Paul Erdős
10. Alverson B, Froese E, Kaplay L, Roweth D Cray® XC™ series network

11. Besta M, Hoefer T (2014) Slim fly: a cost effective low-diameter network topology. In: Proceedings of the International Conference for High Performance Computing, Networking, Storage and Analysis. IEEE Press, pp 348–359
12. Bhatele A, Jain N, Livnat Y, Pascucci V, Bremer P (2016) Analyzing network health and congestion in dragonfly-based supercomputers. In 2016 IEEE International Parallel and Distributed Processing Symposium, IPDPS 2016, Chicago, IL, USA, May 23–27, pp 93–102
13. Biggs N (1993) Algebraic graph theory, Cambridge mathematical library, 2nd edn. Cambridge University Press, Cambridge
14. Bilu Y, Linial N (2006) Lifts, discrepancy and nearly optimal spectral gap. *Combinatorica* 26:495–519
15. Brewer EA, Chong FT, Leighton T (1994) Scalable expanders. In: Proceedings of the Twenty-Sixth Annual ACM symposium on Theory of computing—STOC '94. ACM Press
16. Brouwer AE, Haemers WH (2012) Spectra of graphs. Universitext, Springer, New York
17. Brualdi RA (2011) The mutually beneficial relationship of graphs and matrices, vol 115. American Mathematical Society, Providence
18. Buser P (1982) A note on the isoperimetric constant. *Annales scientifiques de École normale supérieure* 15:213–230
19. Chung F (2016) A generalized Alon-Boppana bound and weak Ramanujan graphs. *Electron J Comb* 23:3–4
20. Chung FRK (1989) Diameters and eigenvalues. *J Am Math Soc* 2:187–187
21. Chung FRK (1997) Spectral graph theory, vol. 92 of CBMS Regional Conference Series in Mathematics, Published for the Conference Board of the Mathematical Sciences, Washington, DC
22. Cioabă SM (2006) Closed walks and eigenvalues of abelian Cayley graphs. *CR Math* 342:635–638
23. Cohen MB (2016) Ramanujan graphs in polynomial time. In: 2016 IEEE 57th Annual Symposium on Foundations of Computer Science (FOCS). IEEE
24. Fiedler M (1973) Algebraic connectivity of graphs. *Czechoslovak Math J* 23:298–305
25. Friedman J (2003) A proof of Alon's second eigenvalue conjecture. In: Proceedings of the Thirty-Fifth ACM Symposium on Theory of Computing—STOC '03. ACM Press
26. Frieze AM (2001) Edge-disjoint paths in expander graphs. *SIAM J Comput* 30:1790–1801
27. Gioiosa R, Tumeo A, Yin J, Warfel T, Haglin D, Betelu S (2017) Exploring data vortex systems for irregular applications. In: 2017 IEEE International Parallel and Distributed Processing Symposium (IPDPS). IEEE
28. Gioiosa R, Warfel T, Yin J, Tumeo A, Haglin D (2016) Exploring data vortex network architectures. In: 2016 IEEE 24th Annual Symposium on High-Performance Interconnects (HOTI). IEEE
29. Gkantsidis C, Mihail M, Saberi A (2003) Conductance and congestion in power law graphs. *SIGMETRICS Perform Eval Rev* 31:148–159
30. Godsil C, Royle G (2001) Algebraic graph theory, vol 207. Graduate Texts in Mathematics. Springer, New York
31. Hawkins C, Small BA, Wills DS, Bergman K (2007) The data vortex, an all optical path multicomputer interconnection network. *IEEE Trans Parallel Distrib Syst* 18:409–420
32. Heydemann M-C (1997) Cayley graphs and interconnection networks. In: *Graph Symmetry*. Springer, Netherlands, pp 167–224
33. Hoory S (2005) A lower bound on the spectral radius of the universal cover of a graph. *J Combin Theory Ser B* 93:33–43
34. Iliadis I, Chrysos N, Minkenberg C (2007) Performance evaluation of the data vortex photonic switch. *IEEE J Sel Areas Commun* 25:20–35
35. Jung-hyun S, Hyeongok L, Moon-suk J (2008) Optimal routing and Hamiltonian cycle in Petersen-torus networks, Busan, South Korea, November 2008. IEEE
36. Kar S, Moura JM (2006) Ramanujan topologies for decision making in sensor networks. In: 44th Allerton Conference on Communication, Control, and Computing, Citeseer
37. Kim J, Dally WJ, Scott S, Abts D (2008) Technology-driven, highly-scalable dragonfly topology. In: 2008 International Symposium on Computer Architecture. IEEE
38. Kleinberg J, Rubinfeld R (1996) Short paths in expander graphs. In 37th Annual Symposium on Foundations of Computer Science (Burlington, VT, 1996). IEEE Computer Society Press, Los Alamitos, CA 1996:86–95
39. Lawler GF, Sokal AD (1988) Bounds on the L^2 spectrum for Markov chains and Markov processes: a generalization of Cheeger's inequality. *Trans Am Math Soc* 309:557–580

40. Leighton FT (2014) Introduction to parallel algorithms and architectures: arrays· trees· hypercubes. Elsevier, Amsterdam
41. Lenzen C, Wattenhofer R (2016) Clex: yet another supercomputer architecture? [arXiv:1607.00298v1](https://arxiv.org/abs/1607.00298v1)
42. Lubotzky A, Phillips R, Sarnak P (1988) Ramanujan graphs. *Combinatorica* 8:261–277
43. Marcus A, Spielman DA, Srivastava N (2013) Interlacing families I: bipartite ramanujan graphs of all degrees. In: 2013 IEEE 54th Annual Symposium on Foundations of Computer Science. IEEE
44. Marcus AW, Spielman DA, Srivastava N (2015) Interlacing families IV: bipartite Ramanujan graphs of all sizes. In: 2015 IEEE 56th Annual Symposium on Foundations of Computer Science. IEEE
45. Margulis GA (1988) Explicit group-theoretical constructions of combinatorial schemes and their application to the design of expanders and concentrators. *Problemy peredachi informatsii* 24:51–60
46. McKay BD, Miller M, Širáň J (1998) A note on large graphs of diameter two and given maximum degree. *J Comb Theory Ser B* 74:110–118
47. Miller M, Širáň J (2005) Moore graphs and beyond: a survey of the degree/diameter problem. *Electron J Comb* 1000:DS14
48. Mohar B (1991) Eigenvalues, diameter, and mean distance in graphs. *Graphs Comb* 7:53–64
49. Mohar B, Alavi Y, Chartrand G, Oellermann O (1991) The laplacian spectrum of graphs. *Graph Theory Comb Appl* 2:12
50. Morgenstern M (1994) Existence and explicit constructions of $q + 1$ regular ramanujan graphs for every prime power q . *J Comb Theory Ser B* 62:44–62
51. Nilli A (1991) On the second eigenvalue of a graph. *Discrete Math* 91:207–210
52. Paturi R, Lu D-T, Ford JE, Esener SC, Lee SH (1991) Parallel algorithms based on expander graphs for optical computing. *Appl Opt* 30:917
53. Pinto EC, Petit C (2018) Better path-finding algorithms in LPS Ramanujan graphs. *J Math Cryptol* 12:191–202
54. Preparata FP, Vuillemin J (1981) The cube-connected cycles: a versatile network for parallel computation. *Commun ACM* 24:300–309
55. Prieto-Castrillo F, Astillero A, Botón-Fernández M (2014) A stochastic process approach to model distributed computing on complex networks. *J Grid Comput* 13:215–232
56. Riess C, Strehl V, Wanka R (2012) The spectral relation between the cube-connected cycles and the shuffle-exchange network. *PARS Parallel Algorithmen Rechnerstrukturen und -Systemsoftware* 29:15–26
57. Sardari NT Complexity of strong approximation on the sphere. [arXiv:1703.02709](https://arxiv.org/abs/1703.02709)
58. Shacham A, Small B, Liboiron-Ladouceur O, Bergman K (2005) A fully implemented 12×12 data vortex optical packet switching interconnection network. *J Lightwave Technol* 23:3066–3075
59. Sinclair A, Jerrum M (1989) Approximate counting, uniform generation and rapidly mixing Markov chains. *Inf Comput* 82:93–133
60. Singla A, Hong C-Y, Popa L, Godfrey PB (2012) Jellyfish: networking data centers randomly. In: Presented as part of the 9th {USENIX} Symposium on Networked Systems Design and Implementation ({NSDI} 12), pp 225–238
61. Tanner RM (1984) Explicit concentrators from generalized N-gons. *SIAM J Algebraic Discrete Methods* 5:287–293
62. Upfal E (1992) An $o(\log n)$ deterministic packet-routing scheme. *J ACM* 39:55–70
63. Valadarsky A, Shahaf G, Dinitz M, Schapira M (2016) Xpander: towards optimal-performance data-centers. In: Proceedings of the 12th International on Conference on emerging Networking Experiments and Technologies—CoNEXT '16. ACM Press
64. Vazirani VV (2001) Approximation algorithms. Springer, Berlin
65. West DB et al (1996) Introduction to graph theory, vol 2. Prentice Hall, Upper Saddle River
66. Yang Q, Bergman K (2002) Performances of the data vortex switch architecture under nonuniform and bursty traffic. *J Lightwave Technol* 20:1242–1247
67. Young SJ Weighted spectrum of the universal cover and an Alon-Boppana result for the normalized Laplacian. Preprint