

Nodal count for a random signing of a graph with disjoint cycles

Lior Alon and Mark Goresky

Abstract. A recent conjecture, inspired by quantum chaos and the Bogomolny–Schmit conjecture, suggests that the nodal count of operators on signed graphs exhibits a universal Gaussian-like behavior. We establish this result for the family of graphs composed of disjoint cycles, which serves as a natural starting point by analogy with quantum graphs. Let G be a simple, connected graph with disjoint cycles, and let h be a real symmetric matrix supported on G (e.g., a discrete Schrödinger operator) that satisfies a certain generic condition. The nodal count $\nu(h, k)$ is defined as the number of edges (i, j) where the k -th eigenvector ϕ changes sign with respect to h , i.e., $h_{ij}\phi_i\phi_j > 0$. We consider the distribution of nodal counts $\nu(h', k)$ over random *signings* h' of h , obtained by changing the sign of some off-diagonal elements. We prove, for each k that $\sigma(h', k) = \nu(h', k) - (k - 1)$ has a binomial distribution $\text{Bin}(\beta, \frac{1}{2})$, where β is the first Betti number of G . Consequently, the conjecture is validated for graphs with disjoint cycles.

1. Introduction

A celebrated theorem of R. Courant [17] states that the k -th eigenfunction for the Dirichlet Laplacian in a planar domain has at most k nodal domains. Over the years, considerable progress has been made in refining, generalizing and applying this result to related questions. U. Smilansky, S. Gnutzmann, and G. Blum [14] proposed that nodal count statistics should predict quantum chaos. They found numerical evidence that the fluctuations of the nodal count, for a chaotic planar domain, obey a universal Gaussian law, a statement that is now known as the Bogomolny–Schmit conjecture [8, 15]. Various models (e.g., percolation) and special cases (e.g. tori and spheres) have been studied but the general conjecture remains wide open.

A more approachable model may be the 1-dimensional case: a quantum graph. A finite quantum graph G with edge lengths ℓ has infinitely many eigenvalues which,

Mathematics Subject Classification 2020: 05C22 (primary); 05C50, 47A10, 57R70, 58K05, 49J52, 57Z05 (secondary).

Keywords: spectral graph theory, nodal count, Morse theory.

for large graphs with incommensurate edge lengths, exhibit the spectral properties of a chaotic system. The k -th eigenfunction has $k - 1 + \sigma(k)$ zeroes, where the *nodal surplus* $\sigma(k)$ is known [9] to be uniformly bounded, $0 \leq \sigma(k) \leq \beta(G)$ by the first Betti number of the graph. In this setting the nodal surplus distribution σ can be rigorously defined [1] as the $N \rightarrow \infty$ limit of the distribution of $\sigma(k)$ when $k \in \{1, \dots, N\}$ is chosen uniformly at random. The universality conjecture for quantum graphs ([2, 19]) states that the nodal surplus distribution σ should approach a Gaussian centered at $\beta(G)/2$, when $\beta(G)$ is large.

Even for quantum graphs, despite considerable numerical evidence in its favor, this conjecture remains open. The nodal surplus distribution is only known in a single situation: when the graph has *disjoint cycles*. In this case, the distribution σ was found [1] to equal the binomial distribution $\text{Bin}(\beta(G), \frac{1}{2})$ – and is independent of edge lengths! – verifying the conjecture for quantum graphs with disjoint cycles.

It is natural to ask the analogous question for discrete Schrödinger operators on finite graphs since the nodal count (see §1.1) of the k -th eigenfunction differs from $k - 1$ by a *nodal surplus* $0 \leq \sigma(k) \leq \beta(G)$, just as in the case of quantum graphs [9]. But a discrete operator on a finite graph has only n (the number of vertices) eigenvalues so the distribution of values $\sigma(k)$ for very large k no longer makes sense. In fact, Alon and Urschel [6] found examples, for arbitrarily large n and β , where the nodal distribution of $\sigma(k) \in \{0, 1, \dots, \beta(G)\}$ for random $k \in \{1, \dots, n\}$ has mean equal to $\beta(G)/n$, which is very far from $\beta(G)/2$ as one might naïvely hope. The lesson for discrete graphs is that randomizing with respect to k (the choice of eigenvalue) does not provide a good model for understanding the distribution of the nodal surplus.

In [4], the present authors randomized an operator on a graph by assigning a random sign¹ to each edge of the graph independently. In [4] it is conjectured that the distribution of $\sigma(k)$ over random signings *and* random k should approach a Gaussian centered at $\beta(G)/2$. However, a surprising recent result [5] proves the distribution of $\sigma(k)$, for random k , and for GOE-random $n \times n$ matrices converges to a semi-circular distribution as $n \rightarrow \infty$. This means the conjecture in [4] is false, since the GOE distribution is invariant under sign changes of the matrix entries. Numerics suggest, nevertheless, that *for each fixed* k , with random signings and large $\beta(G)$, the distribution of values $\sigma(k) \in \{0, 1, \dots, \beta(G)\}$ appears to be Gaussian.

This was verified in [4] in the special case of complete graphs G with (very) dominant on-site potential. In the present paper, inspired by [1], we prove (Theorem 1.3) the same holds for almost all operators supported on a connected graph G , provided

¹The nature of the spectrum of a graph with a random signings was investigated by Y. Bilu and N. Linial [13] and later used by A. Marcus, D. Spielman and N. Srivastava in their resolution [23] of the Kadison–Singer problem.

G has disjoint cycles. In other words, for *any fixed* k , and with random signings, the nodal surplus distribution $\sigma \sim \text{Bin}(\beta(G), \frac{1}{2})$ is binomial (and hence converges to Gaussian for large $\beta(G)$).

Graphs with disjoint cycles do not, to our knowledge, arise naturally. But they may be considered to be the simplest generalization of the fundamental work of Fiedler [18] who analyzed the nodal count on “acyclic” graphs, that is, trees, or graphs with no cycles.

1.1. Graph setting and notations

Let $G = G([n], E)$ be a simple graph on n ordered vertices $[n] := \{1, 2, \dots, n\}$ with a set of edges E and first Betti number $\beta = \text{rank}(H_1(G, \mathbb{R}))$. Write $r \sim s$ if the vertices $r \neq s$ are connected by an edge $(rs) \in E$. An $n \times n$ matrix h is *supported* (resp. strictly supported) on G if for any $r \neq s$, $h_{rs} \neq 0 \implies r \sim s$ (resp. $h_{rs} \neq 0 \iff r \sim s$ for $r \neq s$). Let $\mathcal{S}(G)$ (resp. $\mathcal{A}(G)$, resp. $\mathcal{H}(G)$) denote the vector space of real symmetric (resp. antisymmetric, resp. Hermitian) matrices supported on G . The eigenvalues of a symmetric matrix $h \in \mathcal{S}(G)$ are real and ordered, $\lambda_1(h) \leq \lambda_2(h) \leq \dots \leq \lambda_n(h)$. We say that $\phi \in \mathbb{R}^n$ is *nowhere-vanishing* if $\phi_j \neq 0$ for all j . If ϕ is a nowhere-vanishing eigenvector of h , with simple eigenvalue λ_k , then its *nodal (edge) count* is

$$v(h, k) = |\{(rs) \in E : \phi_r h_{rs} \phi_s > 0\}|.$$

(If $h_{rs} < 0$, as in the case of the graph Laplacian or more generally, a discrete Schrödinger operator, the nodal (edge) count is the number of edges on which ϕ changes sign.) If the graph G is a tree, the nodal count is exactly $v(h, k) = k - 1$ [18], however, this is not the case if G is not a tree [7]. Consequently, the *nodal surplus* for the k -th eigenvalue of h is defined to be

$$\sigma(h, k) := v(h, k) - (k - 1),$$

and by [9] it is bounded: $0 \leq \sigma(h, k) \leq \beta$.

A *signing* of $h \in \mathcal{S}(G)$ is a symmetric matrix h' obtained from h by changing the sign of some of its off-diagonal elements. When considering a random signing h' , we choose an element from the set of $2^{|E|}$ signings uniformly at random. In this way, $\sigma(h', k)$ is a random variable supported on $\{0, 1, \dots, \beta\}$.

A cycle is a path along the graph starting and ending at the same vertex, and it is simple if no other vertex is repeated. We say that G has *disjoint cycles* if distinct simple cycles do not share any vertex. See §3 and Figure 1.




Figure 1. A graph with disjoint cycles.

1.2. Main result

If $\phi \in \mathbb{R}^n$ is an eigenvector of $h \in \mathcal{S}(G)$, in order to avoid double subscripts, we sometimes write $\phi(r) = \phi_r$. To define the nodal count for all signing of $h \in \mathcal{S}(G)$, the matrix h must satisfy the following *generic spectral condition*.

[GSC] The matrix h is strictly supported on G , and every eigenvalue of every signing of h is simple with nowhere vanishing eigenvector.

In Proposition 6.3 we establish that condition [GSC] is indeed generic. The main result of this paper is the following.

Theorem 1.3. *Let G be a simple connected graph with n vertices and disjoint cycles. Suppose $h \in \mathcal{S}(G)$ satisfies [GSC], and let h' be a random signing of h . Then for any $k \in [n]$, the random variable $\sigma(h', k)$ is binomially distributed: the fraction of those signings h' such that $\sigma(h', k) = j$ is $2^{-\beta} \binom{\beta}{j}$. Consequently, as $\beta \rightarrow \infty$, this distribution converges to a Gaussian centered at $\beta/2$ with variance $\beta/4$.*

1.4. Nodal-magnetic relation

Given a graph G with a matrix h as above, the various signings of h lie in a single torus² $\mathbb{T}_h \subset \mathcal{H}(G)$ of so-called *magnetic perturbations* of h . It consists of elements h_α with $(h_\alpha)_{rs} = e^{i\alpha_{rs}} h_{rs}$ with $\alpha \in \mathcal{A}(G)$, see §2.4. We may consider the eigenvalue λ_k to be a sort of Morse function on $\mathbb{T}_h$. A signing $h' \in \mathbb{T}_h$ occurs when all phases are $\alpha_{rs} \in \{0, \pi\}$. These are the “real” points, or “symmetry points” in $\mathcal{M}_h$, the points fixed

²Defined (in [10]) by allowing off-diagonal elements h_{rs} to vary by a phase $e^{i\theta_{rs}} h_{rs}$.

under complex conjugation. It is a theorem of Berkolaiko [10], further explained by Colin de Verdière [16], that each signing $h' \in \mathbb{T}_h$ is a critical point of λ_k , whose Morse index coincides with the nodal surplus for h' . Unfortunately, due to the existence of a group of *gauge transformations* that acts on $\mathbb{T}_h$ and preserves λ_k , each critical point h' is highly degenerate.

The degeneracy in the critical points can be removed by dividing the torus $\mathbb{T}_h$ by the gauge group, §2.5. The result is a torus $\mathcal{M}_h$ whose dimension

$$\beta = |E| - n + 1 = \text{rank}(H_1(G))$$

is the first Betti number of G . The genericity condition [GSC] now implies ([4, Theorem 3.2]) that, for each signing h' of h , the corresponding point $[h'] \in \mathcal{M}_h$ is a non-degenerate critical point of $\lambda_k: \mathcal{M}_h \rightarrow \mathbb{R}$. One might hope that these are the only critical points of λ_k . If this were the case, then we would conclude that λ_k is a *perfect* Morse function, that each critical point contributes to the homology of $\mathcal{M}_h$ in a single degree and hence the nodal surplus is binomially distributed. This situation occurs in [4, Theorem 3.2 and §3.4], where it was proven that the nodal surplus distribution is binomial when G is a complete graph and h has a dominant diagonal.

For generic graphs with a complex cycle structure, experiments indicate the existence of many critical points of λ_k (in addition to the signings of h), which casts doubt on the possibility of using Morse theory to understand the nodal surplus distribution in general.

For generic graphs with disjoint cycles, we believe that each λ_k is a perfect Morse function, but we do not prove it. In [3] it is shown that there are no other *smooth* critical points on $\mathcal{M}_h$, but there may be many points where λ_k has higher multiplicity. The analysis near these points is very complicated, but one expects such points to be topologically regular in the sense of [12]. If this could be established, then it would give a Morse theoretic proof of our main theorem.

1.5. Hypercube toolkit

Instead, we develop a different approach using the combinatorics of the Boolean lattice (§3.6) and two technical steps: (a) the *monotonicity lemma* (Lemma 5.1), adapted from [22], and (b) the *local-global* theorem (Theorem 7.3), adapted from [11]. These results allow us to focus on the one-dimensional trajectories that connect neighboring signings h' and h'' as described in Propositions 3.4 and 3.5. The proof is then outlined in §3.6.

1.6. Bloch variety of periodic operators

The spectral analysis of magnetic operators on finite graphs has an equivalent formulation in terms of periodic graphs, which we briefly review for the sake of completeness, see [22]. The universal abelian cover $\widehat{G}$ of the graph G is a $\mathbb{Z}^\beta$ periodic graph. The matrix h lifts to a $\mathbb{Z}^\beta$ -periodic operator, $\widehat{h}$, on $\widehat{G}$. In Bloch theory, the space $L^2(\widehat{G})$ decomposes into an orthogonal sum (or direct integral) of subspaces $L_\theta^2(\widehat{G})$ corresponding to characters $\theta \in \mathbb{T}^\beta$ of $\mathbb{Z}^\beta$. A natural identification $\mathbb{T}^\beta \cong \mathcal{M}_h$, is described in §2.6 below, say, $\theta \mapsto \alpha$. The operator $\widehat{h}$ restricted to $L_\theta^2(\widehat{G})$ is unitarily equivalent to the magnetic operator h_α . The spectrum of $\widehat{h}$ is therefore the union over the spectrum of h_α for all α .

The graph of the eigenvalues of h_α as functions on $\mathcal{M}_h$ is called the *Bloch variety* (also known as the *dispersion relation manifold*). It is a subset of $\mathcal{M}_h \times \mathbb{R}$ and its projection to $\mathbb{R}$ consists of “bands” corresponding to different k . There is considerable interest in determining the edges of these bands, that is, the maxima and minima of the various $\lambda_k(\widehat{h})$ over the torus $\mathcal{M}_h$.

When translated into this language, a simple consequence of our main theorem is the following. If $\widehat{G}$ is the universal abelian cover of a graph G with disjoint cycles, and h is a generic operator on G , then the edges of the spectral bands for $(\widehat{G}, \widehat{h})$ all arise from the symmetry points in $\mathcal{M}_h$, that is, from the signings of h .

1.7. Probability current

An important ingredient in the proof of Theorem 1.3 is the *probability current* $\mathbb{J}(h, \phi)$ (Definition 4.1), a real anti-symmetric matrix supported on G , which may be interpreted as a gauge invariant divergence-free vector field or as a harmonic 1-form. It is defined for any $h \in \mathcal{H}(G)$ and every eigenvector of h and has a special structure. It vanishes on every bridge³ and is constant on the edges of each simple separated cycle. If the eigenvalue λ is simple and the eigenvector is normalized, then $-2\mathbb{J}$ is the derivative of λ , cf. Proposition 4.2.

2. Recollections on graphs

2.1. As in §1, we consider a simple connected graph G on n ordered vertices numbered $1, 2, \dots, n$. We write $\mathcal{H}_n, \mathcal{S}_n, \mathcal{A}_n$ for the Hermitian, real symmetric, and real antisymmetric $n \times n$ matrices, and we write $\mathcal{H}(G), \mathcal{S}(G), \mathcal{A}(G)$ for those matrices supported on G . If (rs) is an edge in G , write $E[rs]$ for the matrix that is zero except

³A *bridge* is an edge whose removal disconnects the graph.

for $E[rs]_{rs} = 1$ and let $A[rs] = E[rs] - E[sr]$ be the corresponding antisymmetric matrix.

2.2. The space $C_0(G, \mathbb{R})$ of real valued 0-chains consists of (real) linear combinations of vertices. It has a basis $\{e_r\}$ ($1 \leq r \leq n$) with one basis vector per vertex. A 0-chain $a \in C_0(G, \mathbb{R})$ is a formal linear combination of vertices,

$$a = \sum_{r=1}^n a_r e_r.$$

Each edge (rs) of G has a natural orientation (+ or $-$) which is the sign of $s - r$. The space $C_1(G; \mathbb{R})$ of 1-chains consists of formal (real) linear combinations of oriented edges. It has a basis $\{e_{(rs)}\}$ with one element for each edge (rs) with $r < s$. A typical 1-chain is a linear combination

$$\gamma = \sum_{\substack{r \sim s \\ r < s}} \gamma_{rs} e_{(rs)} \quad \text{with } \gamma_{rs} \in \mathbb{R}.$$

The boundary map, $\partial: C_1(G) \rightarrow C_0(G)$ is defined by $\partial e_{(rs)} = e_s - e_r$.

We may consider the space of real antisymmetric matrices $\mathcal{A}(G)$ to be the space of (real valued) 1-forms $\Omega^1(G; \mathbb{R})$, dual to $C_1(G)$ with respect to the bilinear pairing

$$\int_{\gamma} \alpha := \sum_{\substack{r \sim s \\ r < s}} \gamma_{rs} \alpha_{rs}$$

where $\gamma \in C_1(G; \mathbb{R})$ and $\alpha \in \mathcal{A}(G)$.

The space of real valued functions defined on the vertices of G is denoted

$$\Omega^0(G, \mathbb{R}) \cong \mathbb{R}^n$$

and it is naturally identified with the dual space of $C_0(G, \mathbb{R})$. The differential

$$d: \Omega^0(G) \rightarrow \Omega^1(G)$$

is

$$(d\theta)_{rs} = \begin{cases} \theta(s) - \theta(r) & \text{if } r \sim s, \\ 0 & \text{otherwise.} \end{cases}$$

If $\alpha \in \Omega(G, \mathbb{R})$ is a 1-form, its *divergence* is $d^* \alpha \in \Omega^0(G, \mathbb{R})$ where d^* is the adjoint of d with respect to the natural inner products⁴ on $\Omega^0(G)$ and $\Omega^1(G)$, that is,

$$(d^* \alpha)_r = \sum_s \alpha_{rs}.$$

⁴Given by $\langle \theta, \theta' \rangle = \sum_r \theta_r \theta'_r$ and $\langle \alpha, \alpha' \rangle = \sum_{r < s} \alpha_{rs} \alpha'_{rs}$ for $\theta \in \Omega^0(G, \mathbb{R})$ and $\alpha \in \Omega^1(G, \mathbb{R})$.

2.3. Stokes' theorem $\int_{\gamma} d\theta = \int_{\partial\gamma} \theta$ implies that the integration pairing passes to a non-singular dual pairing between the cohomology $H^1(G, \mathbb{R}) = \Omega^1(G)/d\Omega^0(G)$ and the homology $H_1(G, \mathbb{R}) = \ker(\partial)$. Consequently, given $\alpha \in \mathcal{A}(G)$, there exists $\theta \in \Omega^0(G, \mathbb{R})$ such that $\alpha = d\theta$ if and only if $\int_{\gamma} \alpha = 0$ for every cycle γ .

2.4. Action of $\mathcal{A}(G)$

The space $\mathcal{A}_n(\mathbb{R})$ of real $n \times n$ antisymmetric matrices acts on the space $\mathcal{H}_n$ of Hermitian matrices by

$$(\alpha * h)_{rs} = e^{i\alpha_{rs}} h_{rs}$$

with $\alpha' * \alpha * h = (\alpha' + \alpha) * h$. Let $\mathcal{A}_n(2\pi\mathbb{Z})$ be the set of antisymmetric matrices whose entries are integer multiples of 2π . The action factors through the torus $\mathcal{A}_n(\mathbb{R})/\mathcal{A}_n(2\pi\mathbb{Z})$ so that

$$\mathbb{T}(G) = \{\alpha \in \mathcal{A}_n(\mathbb{R})/\mathcal{A}_n(2\pi\mathbb{Z}) : \alpha_{rs} \neq 0 \implies r \sim s\}$$

acts on $\mathcal{H}(G)$. The mapping

$$*: \mathbb{T}(G) \times \mathcal{S}(G) \rightarrow \mathcal{H}(G)$$

is a finite surjective covering. For each $h \in \mathcal{S}(G)$, the orbit

$$\mathbb{T}_h = \mathbb{T}(G) * h$$

is a torus of perturbations⁵ of h . The torus $\mathbb{T}_h$ is preserved under complex conjugation and the fixed points are the intersection $\mathbb{T}_h \cap \mathcal{S}(G)$, which consists of the *signings* of h .

2.5. Gauge equivalence

If $\theta = (\theta_1, \theta_2, \dots, \theta_n) \in \Omega^0(G, \mathbb{R}) \cong \mathbb{R}^n$ and $h \in \mathcal{H}(G)$ then

$$d\theta * h = e^{i\theta} h e^{-i\theta}$$

is conjugate to h , where $e^{i\theta} = \text{diag}(e^{i\theta_1}, e^{i\theta_2}, \dots, e^{i\theta_n})$. Therefore, $\lambda_k(d\theta * h) = \lambda_k(h)$. If $V_{\lambda}(h) = \ker(h - \lambda I)$, then

$$V_{\lambda}(d\theta * h) = e^{i\theta} V_{\lambda}(h). \quad (1)$$

⁵Referred to in [10] as the torus of “magnetic perturbations of h ” because, for the Schrödinger operator, these perturbations arise from the introduction of a magnetic field, cf. [4].

We say the elements h and $h' = d\theta * h$ are *gauge equivalent* and differ by the *gauge transformation* $d\theta$. Geometrically, equation (1) says that eigenvectors ϕ, ϕ' of h and h' differ by changing the phases, $\phi'_r = e^{i\theta_r} \phi_r$. Since their eigenvalues λ_k, λ'_k are equal, the eigenvalue passes to a well-defined function on the set of gauge-equivalence classes of matrices.

We may formally define the gauge group $\mathcal{G} = (\mathbb{R}/2\pi\mathbb{Z})^n$ with action $\theta \diamond h = d\theta * h$, whose orbits are gauge equivalence classes. The quotient of $\mathbb{T}_h$ under gauge equivalence is an abstract torus $\mathcal{M}_h = \mathbb{T}_h // \mathcal{G}$, of dimension β , the *manifold of magnetic perturbations modulo gauge transformations*. We sometimes write $[h] \in \mathcal{M}_h$ for the gauge-equivalence class of h .

Equation (1) reflects an action of the gauge group on vectors $\phi \in \mathbb{C}^n$ with $\theta \diamond \phi = e^{i\theta} \phi$.

2.6. The function Λ_k and choice of basis for $\mathbb{T}^\beta$

Fix a spanning tree in G . Its complement consists of a single edge in each simple cycle. The elements $\alpha \in \mathbb{T}(G)$ that are supported on these edges form a torus $\mathbb{T}^\beta$ that projects isomorphically to the quotient torus $\mathcal{M}_h$. In other words, every element $\alpha * h \in \mathbb{T}_h$ is gauge equivalent to some $\alpha' * h$ where α' is supported on these chosen edges. Thus, $\mathbb{T}^\beta$ is a “lift” to $\mathbb{T}(G)$ of the manifold $\mathcal{M}_h$, as in the following diagram. The composition across the top row is denoted $\Lambda_k: \mathbb{T}^\beta \rightarrow \mathbb{R}$:

$$\begin{array}{ccccccc}
 \mathbb{T}^\beta & \hookrightarrow & \mathbb{T}(G) & \xrightarrow{*h} & \mathbb{T}_h & \hookrightarrow & \mathcal{H}_n \xrightarrow{\lambda_k} \mathbb{R} \\
 & & & & \downarrow & & \nearrow \lambda_k \\
 & & & & \mathcal{M}_h & &
 \end{array}$$

$\cong$

3. Disjoint cycles

3.1. We say a graph G has *disjoint cycles* if distinct simple cycles do not share a vertex, cf. §1.2. Thus, each edge in G is a bridge unless it is contained in a simple cycle. Throughout this section, we fix a graph G with disjoint cycles and a matrix $h \in \mathcal{S}(G)$. We also fix $k \in [n] = \{1, 2, \dots, n\}$ and consider the eigenvalue function λ_k .

3.2. Combinatorics of $\mathbb{T}^\beta$

Choose an ordering of the edges identified in §2.6 (with one edge in each simple cycle). This gives a particular choice of identification

$$(S^1)^\beta \cong \mathbb{T}^\beta \xrightarrow{*h} \mathcal{M}_h. \quad (2)$$

Let $e_1, e_2, \dots, e_\beta \in \mathbb{T}^\beta$ denote the image in $\mathbb{R}^\beta / (2\pi\mathbb{Z})^\beta$ of the standard basis⁶ vectors. Points $\varepsilon = \sum_{i=1}^\beta \varepsilon_i e_i \in \mathbb{T}^\beta$ with coordinates $\varepsilon_i \in \{0, \pi\}$ are called *symmetry points*. By abuse of notation, we write $\varepsilon \in \{0, \pi\}^\beta$. The corresponding matrices $h_\varepsilon = \varepsilon * h$ are the signings of h modulo gauge equivalence.

There are 2^β symmetry points in $\mathbb{T}^\beta$. They form the vertices of a (hyper-)cube

$$\square \subset \mathcal{M}_h$$

whose 1-skeleton consists of edges that connect a symmetry point ε to a neighbor $\varepsilon + \pi e_j \pmod{2\pi}$ (where $j \in [\beta]$). A choice of eigenvalue λ_k determines a partial ordering on the symmetry points,

$$\varepsilon \succeq \varepsilon' \iff \lambda_k(\varepsilon * h) \geq \lambda_k(\varepsilon' * h).$$

For $\varepsilon \in \{0, \pi\}^\beta \subset \mathbb{T}^\beta$, let

$$J_-(\varepsilon) = J_-(\varepsilon, k, h) = \{j \in [\beta] : \lambda_k((\varepsilon + \pi e_j) * h) < \lambda_k(\varepsilon * h)\}.$$

The set $J_-(\varepsilon)$ identifies those neighbors $\varepsilon + \pi e_j$ of ε in the 1-skeleton for which the eigenvalue $\lambda_k(h_\varepsilon)$ decreases.

3.3. Although the proof of our main result (Theorem 1.3) has many technical steps, the ideas are relatively simple, requiring only the following two propositions whose proofs appear in §7. Let G be a simple connected graph with disjoint cycles and suppose $h \in \mathcal{S}(G)$ is generic in the sense of [GSC]. Fix $k \in [n]$ and recall the notation $\Lambda_k(\alpha) = \lambda_k(\alpha * h)$ for $\alpha \in \mathbb{T}^\beta$.

Proposition 3.4. *Each $\varepsilon \in \{0, \pi\}^\beta$ is a non-degenerate critical point of the function $\Lambda_k: \mathbb{T}^\beta \rightarrow \mathbb{R}$. Its Morse index is $\text{ind}(\Lambda_k)(\varepsilon) = |J_-(\varepsilon)|$. The Hessian of the function Λ_k is diagonal with respect to the decomposition (2).*

Proposition 3.5. *The mapping $\{0, \pi\}^\beta \rightarrow \mathcal{P}[\beta]$ (the set of subsets of $[\beta]$), given by $\varepsilon \mapsto J_-(\varepsilon)$ is bijective. This implies that $\{0, \pi\}^\beta$ becomes a Boolean lattice⁷ under the above partial order.*

3.6. Proof of Theorem 1.3

First, we consider the nodal distribution of $\Lambda_k(\varepsilon)$ as ε varies in $\{0, \pi\}^\beta \subset \mathbb{T}^\beta$. By [4, 10, 16] the function Λ_k has a non-degenerate critical point at each $\varepsilon \in \{0, \pi\}^\beta$ and

⁶Each $e_j = A[r_j s_j] = E[r_j s_j] - E[s_j r_j]$ is in fact a matrix in $\mathcal{A}(G)$ defined modulo 2π , and corresponds to one of the particular edges identified in §2.6.

⁷The Boolean lattice on a finite set S is the partially ordered set $\mathcal{P}(S)$ of subsets of S ordered by inclusion.

its Morse index equals the nodal surplus $\sigma(h, k)$ at that point. By Proposition 3.4, this means that the nodal surplus distribution coincides with the distribution of the numbers $|J_-(\varepsilon)|$. Proposition 3.5 implies that the distribution of the numbers $|J_-(\varepsilon)|$, and hence also the nodal surplus distribution for λ_k , is binomial as ε varies in $\{0, \pi\}^\beta$.

Next, we consider the set of signings of h . The set $\{0, \pi\}^\beta * h$ is the quotient of the set of signings of h by the action of the gauge group, or more accurately, the action by a certain subgroup of the gauge group. If $\theta = (\varepsilon_1, \varepsilon_2, \dots, \varepsilon_n) \in \Omega^0(G; \mathbb{R})$ with $\varepsilon_i \in \{0, \pi\}$ and if $h' \in \mathcal{A}(G)$ is a signing of h , then $d\theta * h'$ is another signing. The set of such θ form a group under addition modulo 2π . If h is properly supported on G then this defines a free action of $(\mathbb{Z}/(2))^n$ on the set of signings (cf. [4, §§2.6–2.7]). Each symmetry point $\varepsilon \in \{0, \pi\}^\beta \subset \mathbb{T}^\beta$ corresponds to exactly the same number, $2^{n-\beta}$ of signings. Therefore, the binomial distribution on $\{0, \pi\}^\beta$ becomes the same binomial distribution on the set of signings.

4. Probability current and criticality

Throughout this section, we fix a simple connected graph G with n vertices and $h \in \mathcal{S}(\mathcal{G})$ strictly supported on G .

Definition 4.1. Let $\alpha \in \mathcal{A}(G)$ and set $h_\alpha = \alpha * h$. Given an eigenvector ϕ of h_α , define the *probability current* $\mathbb{J} = \mathbb{J}(h_\alpha, \phi) \in \mathcal{A}(G) = \Omega^1(G, \mathbb{R})$ by

$$\mathbb{J}_{rs} = \Im((h_\alpha)_{rs} \bar{\phi}_r \phi_s) = \Im(e^{i\alpha_{rs}} h_{rs} \bar{\phi}_r \phi_s).$$

We say that the eigenvector ϕ satisfies the *criticality condition* at an edge (rs) if $\mathbb{J}_{rs} = 0$.

We remark that the probability current is defined for any eigenvector whether or not the eigenvalue is simple.

Proposition 4.2. *The probability current $\mathbb{J} = \mathbb{J}(h_\alpha, \phi)$ satisfies the following:*

- (1) $\mathbb{J}$ is gauge-invariant, namely $\mathbb{J}(d\theta * h_\alpha, e^{i\theta} \phi) = \mathbb{J}(h_\alpha, \phi)$;
- (2) $\mathbb{J}$ is divergence free, meaning that $d^* \mathbb{J} = 0$;
- (3) $\mathbb{J}_{rs} = 0$ for every bridge (rs) ;
- (4) $\mathbb{J}$ is constant along the edges of any simple cycle of G that is disjoint from all others;
- (5) if $\lambda(h_\alpha)$, the eigenvalue of ϕ , is simple, then $\mathbb{J}$ is proportional to its derivative,

$$\frac{\partial \lambda(h_\alpha)}{\partial \alpha_{rs}} = \frac{\partial \Lambda}{\partial \alpha_{rs}} = -2 \|\phi\|^2 \mathbb{J}_{rs}.$$

We remark, in particular, if the criticality condition holds on an edge of a disjoint cycle, then it holds on all the edges of that cycle. The proof of Proposition 4.2 will appear after a short review (§4.3) on derivatives of eigenvalues, which is used in the proof.

4.3. Derivatives of eigenvalues

Recall that $A[rs]$ is the antisymmetric matrix with zero entries except for $A[rs]_{rs} = 1$ and $A[rs]_{sr} = -1$. Fix $\alpha \in \mathcal{A}(G)$, and consider the one-parameter family $\alpha(t) = \alpha + tA[rs]$ that goes through α in the (rs) direction. The t -dependence of $\alpha(t) * h$ occurs only in the (rs) and (sr) entries with

$$(\alpha(t) * h)_{rs} = e^{it} e^{i\alpha_{rs}} h_{rs} = e^{it} (h_\alpha)_{rs}.$$

If $\lambda_k(h_\alpha)$ is a simple eigenvalue, then $t \mapsto \lambda_k(\alpha(t) * h)$ is an analytic function of t around $t = 0$, and its derivative at $t = 0$ is the directional derivative of $\lambda_k(\alpha * h)$.

If $\lambda_k(h_\alpha)$ has a non-trivial multiplicity, then the function $\lambda_k(\alpha(t) * h)$ may fail to be differentiable. The theorem of Kato ([21, Theorem 1.8]) and Rellich ([24, Theorem 1]) implies that it is possible to find analytic families of eigenvalues $\mu_k(t) \in \mathbb{R}$ and eigenvectors $\phi_k(t)$, for all $t \in \mathbb{R}$, so that $(\alpha(t) * h)\phi_k(t) = \mu_k(t)\phi_k(t)$. However, the curves $\mu_k(t)$ may cross, when there are multiple eigenvalues, so the index k does not necessarily correspond to the order of these eigenvalues. In other words, as t varies, $\lambda_k(\alpha(t) * h)$ jumps between various analytic branches $\mu_j(\alpha(t) * h)$. Let us choose one such analytic family or “branch,” (μ, ϕ) , and drop the subscript k , and define

$$\Lambda: \mathbb{T}^\beta \rightarrow \mathbb{R}, \quad \Lambda(\alpha') = \mu(\alpha' * h).$$

Using Leibniz’ dot notation to denote derivative with respect to t , and differentiating $h(t)\phi = \mu(t)\phi(t)$ gives

$$(\dot{h}(t) - \dot{\mu}(t))\phi(t) + (h(t) - \mu(t))\dot{\phi}(t) = 0 \quad (3)$$

As in [11, Lemma 2.5] or [4, §5.2], taking the inner product with ϕ where $\|\phi\| = 1$, using that h is Hermitian, and evaluating at $t = 0$ gives the directional derivative of the eigenvalue μ along this branch:

$$\begin{aligned} \frac{\partial \Lambda}{\partial \alpha_{rs}}(\alpha) &= \dot{\mu} = \frac{d}{dt} \mu(\alpha + t(A[rs]) * h) \Big|_{t=0} = \langle \phi, \dot{h}\phi \rangle \\ &= i(\bar{\phi}_r \phi_s (h_\alpha)_{rs} - \bar{\phi}_s \phi_r (\bar{h}_\alpha)_{rs}) = -2\Im((h_\alpha)_{rs} \bar{\phi}_r \phi_s) \end{aligned} \quad (4)$$

where $\phi_r = \phi(r)$ denotes the value of ϕ on the vertex (r) .

For later applications in equation (12), consider the case when $\frac{\partial \Lambda}{\partial \alpha_{rs}}(\alpha) = 0$. We can differentiate (3) once again to obtain, as in [11, Lemma 2.6],

$$\langle \phi, \ddot{h}\phi \rangle = -\Re((h_\alpha)_{rs} \bar{\phi}_r \phi_s) = -((h_\alpha)_{rs} \bar{\phi}_r \phi_s) \quad (5)$$

and

$$\frac{\partial^2 \Lambda}{\partial \alpha_{rs}^2} = \ddot{\mu} = \langle \phi, \ddot{h}\phi \rangle + 2\Re(\langle \phi, \dot{h}\phi \rangle). \quad (6)$$

4.4. Proof of Proposition 4.2

The gauge invariance, $\mathbb{J}(e^{i\theta} h_\alpha e^{-i\theta}, e^{i\theta} \phi) = \mathbb{J}(h_\alpha, \phi)$, is straightforward from the definition. The divergence is

$$\begin{aligned} (d^* \mathbb{J})_r &= \sum_s \Im(\bar{\phi}_r (h_\alpha)_{rs} \phi_s) = \Im\left(\bar{\phi}_r \sum_s (h_\alpha)_{rs} \phi_s\right) \\ &= \Im(\bar{\phi}_r \lambda \phi_r) = \lambda \Im(|\phi_r|^2) = 0. \end{aligned}$$

If removing an edge $E = (rs)$ separates the graph into two pieces, say G_A and G_B , let $\theta \in \Omega^0(G)$ take the value 1 on G_B and 0 on G_A . Then $d\theta$ is supported on E and

$$\mathbb{J}_{rs} = \langle d\theta, \mathbb{J} \rangle = \langle \theta, d^* \mathbb{J} \rangle = 0.$$

Similarly, if E, E' are two edges in a simple cycle that is disjoint from all others, then removing both separates the graph into two pieces. Taking θ as above,

$$\mathbb{J}(E) - \mathbb{J}(E') = \langle d\theta, \mathbb{J} \rangle = 0.$$

Part (5) is a restatement of equation (4).

Lemma 4.5 (Partial criticality). *Let $\alpha \in \mathcal{A}(G)$ and set $h_\alpha = \alpha * h$. Let ϕ be an eigenvector of simple eigenvalue of h_α and let $\mathbb{J} = \mathbb{J}(h_\alpha, \phi)$ be the probability current. Suppose there is a bridge that splits the graph G into G_A and G_B . If h_α is real on the $G_B \times G_B$ block, then $\mathbb{J}$ vanishes on that block:*

$$h_\alpha|_{G_B} \in \mathcal{S}(G_B) \implies \mathbb{J}|_{G_B} = 0.$$

Proof. Let (rs) denote the bridge with $s \in G_A$ and $r \in G_B$. By changing gauge and scaling h_α if needed, we can assume that $(h_\alpha)_{rs} = 1$. Let e_s and e_r be the corresponding standard basis vectors. In the block decomposition to G_A, G_B , we write $h_\alpha = A \oplus B + e_r e_s^* + e_s e_r^*$.

Suppose the simple eigenvalue of interest is $\lambda = 0$ (otherwise replace h_α with $h_\alpha - \lambda I$), and let $\phi = (\phi_A, \phi_B)$ denote its normalized eigenvector. We need to show that if B is real, then ϕ_B is (proportional to) a real vector, in which case $\mathbb{J}|_{G_B} = 0$.

If $\phi_A = 0$, then $\phi_B \in \ker(B)$ and we are done. So, assume $\phi_A \neq 0$. By Proposition 4.2 (3), we know that $\text{Im}[(h_\alpha)_{rs}\bar{\phi}_r\phi_s] = \text{Im}[\bar{\phi}_r\phi_s] = 0$. By scaling ϕ if needed, we can therefore assume that $\phi(r)$ and $\phi(s)$ are real.

We will now show that $\phi' := (\phi_A, \bar{\phi}_B)$ is also in $\ker(h_\alpha)$. Since the kernel is one-dimensional and $\phi_A \neq 0$, this will mean that $\phi' = \phi$ so $\phi_B = \bar{\phi}_B$. Calculate

$$h_\alpha\phi = (A\phi_A + \phi(r)e_s, B\phi_B + \phi(s)e_r) = (0, 0).$$

But $\phi(r) = \phi'(r)$, $\phi(s) = \phi'(s)$, so

$$\begin{aligned} h_\alpha\phi' &= (A \oplus B + e_re_s^* + e_se_r^*)\phi' \\ &= (A\phi_A + \phi(r)e_s, B\bar{\phi}_B + \phi(s)e_r) \\ &= (0, B\bar{\phi}_B + \phi(s)e_r). \end{aligned}$$

Since B and $\phi(s)$ are real, $B\bar{\phi}_B + \phi(s)e_r = \overline{(B\phi_B + \phi(s)e_r)} = 0$. ■

We now return to the special case of G that has disjoint cycles. Recall that each $\varepsilon \in \{0, \pi\}^\beta$ is a non-degenerate critical point of $\Lambda_k: \mathbb{T}^\beta \rightarrow \mathbb{R}$.

Corollary 4.6. *Suppose G has disjoint cycles and $h \in \mathcal{S}(G)$ satisfies [GSC]. Then for each k , the Hessian of Λ_k at any $\varepsilon \in \{0, \pi\}^\beta \subset \mathbb{T}^\beta$ is diagonal with respect to the basis of $\mathbb{T}^\beta$ that was chosen in §2.6.*

Proof. Fix k and ε . We work in the previously chosen (§2.6) basis of $T_\varepsilon\mathbb{T}^\beta \cong \mathbb{R}^\beta$ given by the choice of a single edge per cycle of G , say $(r_j, s_j) \in \gamma_j$. We will show that

$$\frac{\partial^2 \Lambda_k}{\partial \alpha_{r_1 s_1} \partial \alpha_{r_2 s_2}}(\varepsilon) = 0. \quad (7)$$

(All other off-diagonal terms vanish for the same reason.)

Since the cycles are disjoint, there exists a bridge that separates the graph into two parts, G_A, G_B with $\gamma_1 \subset G_A$ and $\gamma_2 \subset G_B$. Let $\alpha(t) = \varepsilon + tA[r_1, s_1]$ and let $h_t = \alpha(t) * h$. The matrix h_t is real except for the (r_1, s_1) and (s_1, r_1) entries so we may apply Lemma 4.5 and Proposition 4.2 (5) to conclude that

$$\frac{\partial \Lambda_k}{\partial \alpha_{r_2, s_2}}(\alpha(t)) = 0$$

for all t around 0. Differentiating with respect to t at $t = 0$ gives equation (7). ■

5. Monotonicity

Lemma 5.1 (Monotonicity). *Suppose G has a cycle γ disjoint from all others and let $h \in \mathcal{S}(G)$ that satisfies [GSC]. Consider any one-parameter family $h_t = \alpha_t * h$ with α_t supported on γ and $\int_\gamma \alpha_t = t$ for all $t \in [0, \pi]$. Then, $t \mapsto \lambda_k(h_t)$ is strictly monotone in $t \in [0, \pi]$ for all $k \in [n]$.*

We remark that, up to gauge equivalence, we may suppose that α is supported on a single edge of γ . This means the family h_t traverses a single segment in the 1-skeleton of the hypercube $\square \subset \mathcal{M}_h$ from §3.2. The monotonicity lemma *only* applies to these special paths. The rest of §5 is devoted to the proof Lemma 5.1, which appears finally in §5.4.

Lemma 5.2 (Flat band criteria). *Suppose G has a cycle γ disjoint from all others, and (12) is an edge in γ . Let $h \in \mathcal{S}(G)$ and consider a one-parameter family $h_t = \alpha_t * h$ where $\alpha_t \in \mathcal{A}(G)$ satisfies $\alpha_t = \alpha_0$ outside of γ and $\int_\gamma \alpha_t = t$ for all $t \in [0, \pi]$. Suppose there exists $t_0 \in (0, \pi)$, and an eigenvector ϕ of h_{t_0} with eigenvalue λ , such that $\mathbb{J}(h_{t_0}, \phi)_{12} = 0$. Then λ is a common eigenvalue of all h_t with $t \in [0, \pi]$.*

Proof. Without loss of generality, assume that $\lambda = 0$, so that $h_{t_0}\phi = 0$. We need to provide a family of vectors ϕ_t such that $h_t\phi_t = 0$ for all t . We will show that $\mathbb{J}_{12} = 0$ implies that

- (i) either there is an edge (rs) in γ such that $\phi(r) = 0$ and $\phi(s) = 0$,
- (ii) or there is a vertex r in γ such that $\phi(r) = 0$ and $\deg(r) \geq 3$.

We will also show that each of these conditions is sufficient for constructing ϕ_t such that $h_t\phi_t = 0$ for all t . To ease notation, let $\alpha = \alpha_{t_0}$. To avoid triple subscripts, write $\alpha(rs)$ for α_{rs} .

First, we show that (i) is sufficient. Let (rs) be an edge in γ such that $\phi(r) = \phi(s) = 0$. Up to gauge equivalence, we may assume that $\alpha_t = \alpha$ on $G \setminus \gamma$, that $\alpha_t(rs) = t$, and that α_t vanishes on all the other edges in γ . Then $h_t\phi = h_{t_0}\phi = 0$ for all t so we may take $\phi_t = \phi$. Next, we show, using $\mathbb{J}_{12} = 0$, that if (i) fails, then (ii) must hold. Assume (i) fails, namely,

- (A) for every edge (rs) in γ , $\phi(r)$ and $\phi(s)$ are not both zero.

By Proposition 4.2, $\mathbb{J} = \mathbb{J}(h_{t_0}, \phi)$ is gauge invariant and constant on γ , so we have $\mathbb{J}_{rs} = \mathbb{J}_{12} = 0$ for every edge (rs) in γ . To prove that (ii) holds, up to change of gauge, we may assume that ϕ is real, cf. equation (1). In this case,

$$\mathbb{J}_{rs} = \phi(r)\phi(s)h_{rs}\sin(\alpha(rs)) = 0 \quad \text{for any } (rs) \in \gamma,$$

so $\alpha(rs) = 0 \pmod{\pi}$ when $\phi(r)\phi(s) \neq 0$, and so (A) gives

$$\sum_{\phi(r)=0} \alpha(sr) + \alpha(rt) = \int_{\gamma} \alpha_{t_0} = t_0 \neq 0 \pmod{\pi}, \quad (8)$$

where $s < r < t$ denote the neighbors of r . Now, suppose r is a vertex in γ of degree 2 with $\phi(r) = 0$. Let $s < r < t$ be the two vertices attached to r . Then $\phi(s) \neq 0$ and $\phi(t) \neq 0$ by (A) so $(h_{t_0}\phi)_r = 0$ reads

$$\phi(s)h_{sr}e^{i\alpha(sr)} + 0 + \phi(t)h_{rt}e^{i\alpha(rt)} = 0$$

which implies

$$\alpha(sr) + \alpha(rt) = 0 \pmod{\pi}$$

whenever $\deg(r) = 2$ and $\phi(r) = 0$. Adding over all vertices r of degree 2 such that $\phi(r) = 0$ gives

$$\sum_{\substack{\phi(r)=0 \\ \deg(r)=2}} \alpha(sr) + \alpha(rt) = 0 \pmod{\pi} \quad (9)$$

where, as before, $s < r < t$ denote the neighbors of r . The terms in this sum are disjoint by (A). Since the sums in (9) and (8) are not equal, then there must be a vertex $r \in \gamma$ with $\deg(r) \neq 2$ and $\phi(r) = 0$, so (ii) holds.

Finally, assuming (ii), we construct ϕ_t . Without loss of generality, we may suppose that we have consecutive vertices $1 < 2 < 3$ in γ with $\phi(2) = 0$, $\deg(2) \geq 3$. Up to gauge equivalence, we may assume that $(\alpha_t)_{12} = t$, $(\alpha_t)_{rs} = 0$ for all other edges (rs) in γ , and $\alpha_t = \alpha$ on $G \setminus \gamma$.

Let H denote the union of connected components of $G \setminus \gamma$ that are connected to vertex 2 in G . We will show there exists $c(t) \in \mathbb{C}$ with $c(t_0) = 1$ so that the vector ϕ_t defined by

$$(\phi_t)_r = \begin{cases} c(t)\phi_r & \text{for } r \in H, \\ \phi_r & \text{for } r \notin H, \end{cases}$$

satisfies $h_t\phi_t = 0$ for all t . By our assumption on G , the only vertex in $G \setminus H$ with a neighbor in H is the vertex 2 on which $\phi(2) = 0$. Therefore, for any t and any edge (r, s) ,

$$(h_t)_{rs}\phi_t(s) = \begin{cases} (h_{t_0})_{rs}\phi(s) & \text{for } r \in H, r \neq 2, \\ C(t)(h_{t_0})_{rs}\phi(s) & \text{for } r \in G \setminus H, \end{cases}$$

from which we conclude that $(h_t\phi_t)_r \propto (h_{t_0}\phi)_r = 0$ for all $r \neq 2$. To understand the situation at vertex 2, let

$$F_2(\alpha, \phi) = \sum_{\substack{r \in H \\ r \sim 2}} h_{2r} e^{i\alpha(2r)} \phi(r).$$

Then $F_2(\alpha_t, \phi_t) = c(t)F_2(\alpha, \phi)$ and at vertex 2 we have

$$(h_{i_0}\phi)_2 = \phi(1)h_{12}e^{it_0} + \phi(3)h_{23} + F_2(\alpha, \phi) = 0.$$

Therefore,

$$-F_2(\alpha, \phi) = \phi(1)h_{12}e^{it_0} + \phi(3)h_{23} \neq 0,$$

since e^{it_0} is not real. The eigenvalue equation at vertex 2 becomes

$$\begin{aligned} (h_t\phi_t)_2 &= \phi(1)h_{12}e^{it} + \phi(3)h_{23} + F_2(\alpha_t, \phi_c) \\ &= \phi(1)h_{12}e^{it} + \phi(3)h_{23} + c(t)F_2(\alpha, \phi). \end{aligned}$$

It is left to choose $c(t)$ so that $(h_t\phi_t)_2 = 0$, namely

$$c(t) = -\frac{\phi(1)h_{12}e^{it} + \phi(3)h_{23}}{F_2(\alpha, \phi)} = \frac{\phi(1)h_{12}e^{it} + \phi(3)h_{23}}{\phi(1)h_{12}e^{it_0} + \phi(3)h_{23}}. \quad \blacksquare$$

Lemma 5.3. *In the setting of Lemma 5.1, if (rs) is an edge in γ and ϕ is a normalized eigenvector of h_t , for some $t \in (0, \pi)$, then $\mathbb{J} = \mathbb{J}(h_t, \phi)$ has $\mathbb{J}_{rs} \neq 0$.*

In particular, if ϕ and ϕ' are eigenvectors of the same eigenvalue of h_t , then $\mathbb{J}(h_t, \phi)_{rs}$ and $\mathbb{J}(h_t, \phi')_{rs}$ share the same sign.

Proof. Since h satisfies [GSC], then each of the eigenvalues has $\Lambda_k(\alpha) = \lambda_k(\alpha * h)$ has a non-degenerate critical point at $\alpha = 0$, namely at $h = h_0$, whose Hessian is diagonal by Corollary 4.6. In particular, for any $k \in [n]$, $\Lambda_k(\alpha_t) = \lambda_k(h_t)$ is not constant around $t = 0$. This means that $\mathbb{J}_{rs} \neq 0$ for any normalized eigenvector of any h_t with $t \in (0, \pi)$, otherwise we would get a “flat band,” namely a constant eigenvalue $\lambda_k(h_t) \equiv \lambda$ for all t around $t = 0$ by Lemma 5.2. This concludes the first part.

Now, let $V = \ker(h_t - \lambda_k(h_t))$ be some eigenspace of some h_t with $t \in (0, \pi)$, and assume $\dim(V) \geq 2$. Then the map $\phi \mapsto \mathbb{J}(h_t, \phi)$ is a continuous map from $V \setminus \{0\}$ (which is connected) to $\mathbb{R} \setminus \{0\}$ so its image must lie either in $\mathbb{R}_{>0}$ or in $\mathbb{R}_{<0}$. $\blacksquare$

5.4. Proof of Lemma 5.1

The statement is gauge invariant, so we may fix the gauge such that α is supported on a single edge, say, (12). By Kato [21, Theorem 1.8] or Rellich [24, Theorem 1], or Wimmer [25], since this is a one-parameter analytic family of Hermitian matrices, the ordered eigenvalues $(\lambda_1 \leq \dots \leq \lambda_n)$ and eigenvectors $(\phi_1, \dots, \phi_n)$ of h extend analytically to eigenvalues and normalized eigenvectors $(\mu_k(t), \phi_k(t))_{k=1}^n$ of h_t , although a priori their order may not be preserved. The derivative

$$\dot{\mu}_k(t) = \frac{d}{dt}\mu_k(t) = \langle \phi_k(t), \dot{h}_t \phi_k(t) \rangle = -2\mathbb{J}(h_t, \phi_k(t))_{12} \quad (10)$$

was calculated in (4). Since $\mathbb{J}(h_t, \phi_k(t))_{12} \neq 0$ for all k and all $t \in (0, \pi)$ by Lemma 5.3, then each $\mu_k(t)$ is strictly monotone in $t \in [0, \pi]$. If all eigenvalues are simple, this proves that $\lambda_k(h_t) = \mu_k(t)$ is monotone for $t \in [0, \pi]$.

If the eigenvalue has a non-trivial multiplicity, say $\mu_k(t) = \mu_{k'}(t)$, then it suffices to know that the derivatives $-2\mathbb{J}(h_t, \phi_k(t))_{12}$ and $-2\mathbb{J}(h_t, \phi_{k'}(t))_{12}$ have the same signs. The second part of Lemma 5.3 ensures this is the case. ■

5.5. Remark

Lemma 5.3 and equation (10) mean that the restriction of the Hermitian form $\dot{h}_t$ to the eigenspace of h_t is sign-definite, which is exactly the condition of [12] for a point of multiplicity to be topologically regular (the BZ condition), see Appendix B.

6. Genericity

6.1. The purpose of this section is to show that the conditions [GSC] of §1.2 are indeed generic. The first (surprisingly tricky) step is to prove the existence of [GSC] matrices on a tree.

Lemma 6.2. *If a graph G of n vertices is a connected tree, there exists a strictly supported matrix $h \in \mathcal{S}(G)$ with simple spectrum and non-vanishing eigenvectors.*

Proof. For $n = 2$ take $h = \begin{pmatrix} 0 & 1 \\ 1 & 0 \end{pmatrix}$. Now, assume the lemma holds for graphs with $\leq n - 1$ vertices. Consider a connected graph G with n vertices. Without loss of generality, we may assume that vertex $\{n\}$ is a leaf and that $(n - 1, n)$ is an edge of G . Let G' be the tree obtained by removing vertex n from G . By induction, there exists $A \in \mathcal{S}(G')$ that is strictly supported, with simple eigenvalues $\lambda_1(A), \dots, \lambda_{n-1}(A)$ and non-vanishing eigenvectors $\psi_1, \dots, \psi_{n-1}$. By changing $A \mapsto A - cI_{n-1}$ and choosing c to avoid⁸ finitely many values, we may assume that $\lambda_j < 0$ for all j , and that

$$u = A^{-1}e_{n-1} = \sum_{j=1}^{n-1} \lambda_j^{-1} \langle \psi_j, e_{n-1} \rangle \psi_j,$$

is nowhere vanishing on G' , where $e_{n-1} \in \mathbb{R}^{n-1}$ is the standard unit vector. Define the analytic one-parameter family with $t \in \mathbb{R}$,

$$h_t = \begin{pmatrix} A & t e_{n-1} \\ t e_{n-1}^T & 0 \end{pmatrix} \in \mathcal{S}(G).$$

⁸For each s , with $1 \leq s \leq n - 1$, $c \neq \lambda_s$ and c is not a root of the non-zero polynomial $c \mapsto \sum_{j=1}^{n-1} \psi_j(s) \langle \psi_j, e_{n-1} \rangle \prod_{t \neq j} (\lambda_t - c)$.

The eigenvalues of h_0 are simple and are equal to $\lambda_j(h_0) = \lambda_j(A)$ for $j \leq n-1$ and $\lambda_n(h_0) = 0$, with eigenvectors $\begin{pmatrix} \psi_0^1 \\ 0 \end{pmatrix}, \dots, \begin{pmatrix} \psi_0^{n-1} \\ 0 \end{pmatrix}$ and e_n . By the Kato–Rellich theorem ([21, 24, 25]), these extend analytically to eigenvalues and eigenvectors of h_t . For sufficiently small $t > 0$, the spectrum will remain simple, and the first $n-1$ eigenvectors of h_t will remain non-zero on the vertices of G' . All eigenvectors of h_t are non-zero on the last vertex, $\{n\}$, since G is a tree, the spectrum is simple, and the vertex $\{n\}$ is not of degree two [18, Corollary 2.7]. The n -th eigenvector of h_t has the form

$$v = e_n + tv' + O(t^2).$$

We claim $v' = -\begin{pmatrix} u \\ 0 \end{pmatrix}$ (so does not vanish on any vertices of G'). In fact,

$$h_t(e_n + tv') = h_t \begin{pmatrix} -tA^{-1}e_{n-1} \\ 1 \end{pmatrix} = \begin{pmatrix} 0 \\ t^2(A^{-1})_{n-1, n-1} \end{pmatrix} = O(t^2). \quad \blacksquare$$

Proposition 6.3. *Let $G = G([n], E)$ be a finite simple connected graph. The set of matrices*

$$\mathcal{O} = \{h \in \mathcal{S}(G) : h \text{ satisfies [GSC]}\}$$

is open and dense in $\mathcal{S}(G)$. Its complement is contained in a closed semi-algebraic⁹ subset of $\mathcal{S}(G)$ of codimension ≥ 1 .

Proof. When eigenvalues are simple, the eigenvalues and eigenvectors vary continuously with the matrix, so the set of matrices satisfying [GSC] is open. If a matrix $h' \in \mathcal{S}(G)$ fails to satisfy [GSC] then there is a signing $\varepsilon \in \{0, \pi\}^E$ for which $h = \varepsilon * h'$ lies in at least one of the following sets:

- (i) the set of matrices in $\mathcal{S}(G)$ that are not strictly supported on G , that is, $h_{ij} = 0$ for some $(ij) \in E$,
- (ii) the set of matrices in $\mathcal{S}(G)$ that have a multiple eigenvalue,

$$\text{discriminant}(h) = 0,$$

or

- (iii) the set of matrices in $\mathcal{S}(G)$ that have a simple eigenvalue with an eigenvector that vanishes at some vertex.

The sets (i) and (ii) are zero sets of polynomials that are not the zero polynomial¹⁰ on $\mathcal{S}(G)$, so these are algebraic subsets of $\mathcal{S}(G)$ with positive codimension in $\mathcal{S}(G)$.

⁹A semi-algebraic subset of a real vector space is a finite Boolean combination of sets defined by polynomial equalities $f(x) = 0$ and inequalities $f(x) < 0$.

¹⁰To see that $\text{discriminant}(h) \neq 0$ for some $h \in \mathcal{S}(G)$, take $h = \text{diag}(1, 2, \dots, n)$.

Since the class of semi-algebraic subsets is preserved by any change of signing, it suffices to show that the set (iii) is semi-algebraic with positive codimension in $\mathcal{S}(G)$.

Recall (see for example, [20, §0.8.2]) that if λ is a simple eigenvalue of h with eigenvector v , then the *adjugate* is defined as

$$A := \text{adj}(h - \lambda.I) = cvv^*$$

where $c \in \mathbb{C}$ is non-zero. If v does not vanish at any vertex, then $A_{ij} = cv_i v_j^* \neq 0$ for all (i, j) . This gives a collection of algebraic subsets $Z'_{ij} \subset \mathcal{S}(G) \times \mathbb{R}$,

$$Z'_{ij} = \{(h, \lambda) \in \mathcal{S}(G) \times \mathbb{R} : \det(h - \lambda.I) = 0 \text{ and } (\text{adj}(h - \lambda.I))_{ij} = 0\}.$$

For a given h , only finitely many λ occur, so this is an algebraic subset of $\mathcal{S}(G) \times \mathbb{R}$ and its projection Z_{ij} to $\mathcal{S}(G)$ is semi-algebraic by the Tarski–Seidenberg theorem. We claim that the union

$$Z := \bigcup_{i,j} Z_{ij} \subset \mathcal{S}(G)$$

has codimension ≥ 1 . This will complete the proof of the theorem, for if an element $h \in \mathcal{S}(G)$ does not lie in this union, and if it has simple eigenvalues, then the corresponding eigenvectors do not vanish on any vertex.

Let $h \in Z$. Choose a spanning tree in G and apply Lemma 6.2 to find a matrix $B \in \mathcal{S}(G)$ with simple eigenvalues and nowhere vanishing eigenvectors. Consider the family $h_t = (1 - t)h + tB$ where $t \in \mathbb{R}$. We claim, for sufficiently small $t > 0$, that h_t also has simple eigenvalues and nowhere vanishing eigenvectors. This will prove that Z does not contain an open neighborhood of h so its codimension at h is ≥ 1 .

According to the theorem of Kato and Rellich [21, 24, 25], it is possible to find eigenvalues $\lambda_i(t)$ and eigenvectors $\psi_i(t)$ ($1 \leq i \leq n$) of $h_t = (1 - t)h + tB$ that vary analytically with $t \in \mathbb{R}$. Therefore, the numbers $\lambda_i(t) - \lambda_j(t)$ vary analytically with t and the values of $\psi_i(t)$ on each vertex also vary analytically. For $t = 1$, the eigenvalues of h_t are distinct and the eigenvectors are nowhere vanishing.

It follows that there is a discrete set of values $t \in \mathbb{R}$ for which h_t has a multiple eigenvalue and there is a discrete set of values $t \in \mathbb{R}$ such that h_t has an eigenvector that vanishes at one or more vertices. Consequently, for $t > 0$ sufficiently small (so as to miss these discrete sets), the matrix h_t will have distinct eigenvalues and nowhere vanishing eigenvectors, as claimed. ■

6.4. Let G be a finite graph. Let $\mathcal{B}$ denote the set of matrices $h \in \mathcal{S}(G)$ that satisfy

(✱) any two gauge-inequivalent signings $\varepsilon * h, \varepsilon' * h$ have distinct eigenvalues.

Lemma 6.5. *The set $\mathcal{B}$ is open and dense in $\mathcal{S}(G)$ and its complement is contained in an algebraic subset of codimension ≥ 1 .*

Proof. We may assume that G is connected. If G is a tree and $h \in \mathcal{S}(G)$, then every signing of h is gauge equivalence to h so we may assume $\beta(G) \geq 1$. First, consider the case that $\beta(G) = 1$ so that G contains a unique cycle. Fix an edge (rs) in this cycle. For any $h \in \mathcal{S}(G)$, there is only one gauge-equivalence class of signings $\varepsilon * h$ of h and it corresponds to changing the sign of h_{rs} (and of h_{sr}).

Let $Q_\varepsilon(h)$ denote the discriminant of the $2n \times 2n$ matrix $(h) \oplus (\varepsilon * h)$. The set $Q_\varepsilon^{-1}(0)$ is an algebraic subset of $\mathcal{S}(G)$ which contains the complement of $\mathcal{B}$. If $Q_\varepsilon^{-1}(0)$ contains an open subset of $\mathcal{S}(G)$, then it is all of $\mathcal{S}(G)$; otherwise it has codimension ≥ 1 . We will assume that $Q_\varepsilon(h) = 0$ for all $h \in \mathcal{S}(G)$ and arrive at a contradiction.

In this one-dimensional case, the hypercube of §3.2 is just an interval whose endpoints are h and $\varepsilon * h$. Let $V = \text{diag}(1, 2, \dots, n)$. Let $\xi \in \mathcal{S}(G)$ (strictly supported on G) sufficiently small such that $h := V + \xi \in \mathcal{O}$ and $\varepsilon * h \in \mathcal{O}$ (such ξ exists by Proposition 6.3). The eigenvalues of h are distinct; the eigenvalues of $\varepsilon * h$ are distinct. Therefore, if $Q_\varepsilon(h) = 0$, then h and $\varepsilon * h$ share an eigenvalue, say, $\lambda_k(h) = \lambda_{k'}(\varepsilon * h)$. If ξ is sufficiently small, the eigenvalues of h and of $\varepsilon * h$ are small perturbations of the eigenvalues of V , which are distinct integers, hence $k = k'$. But this contradicts the monotonicity Lemma 5.1.

We conclude that for any graph G with $\beta(G) = 1$, the function Q_ε vanishes identically on $\mathcal{S}(G)$. Now, consider the case of a general graph $\beta(G) \geq 1$. For a general signing $\varepsilon, \varepsilon' \in \{0, \pi\}^\beta$, set

$$Q_{\varepsilon, \varepsilon'}(h) = \text{discr}((\varepsilon * h) \oplus (\varepsilon' * h)).$$

The complement of $\mathcal{B}$ is contained in the algebraic set

$$Z := \bigcup_{\substack{\varepsilon, \varepsilon' \in \{0, \pi\}^\beta \\ \varepsilon \neq \varepsilon'}} Q_{\varepsilon, \varepsilon'}^{-1}(0) = \left(\prod_{\substack{\varepsilon, \varepsilon' \in \{0, \pi\}^\beta \\ \varepsilon \neq \varepsilon'}} Q_{\varepsilon, \varepsilon'} \right)^{-1}(0).$$

The set Z is a finite union of sets of the form $Q_\varepsilon^{-1}(0)$. To see that each of these sets has codimension ≥ 1 , suppose otherwise. Then there exists a signing ε so that $Q_\varepsilon(h) = 0$ for all $h \in \mathcal{S}(G)$.

Choose a spanning tree in G . Label the edges $e_1, e_2, \dots, e_\beta$ in the complement and express $\varepsilon = \sum \varepsilon_i e_i$ as in §2.6 and §3.2. Arrange the labeling so that $\varepsilon_1 \neq 0$. The graph G' obtained from G by removing the edges $e_2, e_3, \dots, e_\beta$ has $\beta(G') = 1$. The signing ε of G becomes a signing $\eta = \varepsilon_1$ on G' , that is, a change of sign on the remaining edge e_1 . Moreover, any $h' \in \mathcal{S}(G')$ can be obtained as a limit of $h \in \mathcal{S}(G)$ by allowing $h_{rs} \rightarrow 0$ where (rs) varies over the edges $e_2, e_3, \dots, e_\beta$. Since $Q_\varepsilon(h)$ is a continuous function of h , it vanishes on this limiting value, h' . This proves that $Q_\eta(h') = 0$ for all $h' \in \mathcal{S}(G')$ which contradicts the conclusion from the first paragraph. ■

7. Proofs of Propositions 3.4 and 3.5

7.1. Proof of Proposition 3.4

Recall from §3.6 that G is a simple connected graph with disjoint cycles, $h \in \mathcal{S}(G)$ is generic in the sense of [GSC], and $\Lambda_k: \mathbb{T}^\beta \rightarrow \mathbb{R}$ is $\Lambda_k(\alpha) = \lambda_k(\alpha * h)$.

It was shown in [4, 10, 16] that each $\varepsilon \in \{0, \pi\}^\beta \subset \mathbb{T}^\beta$ is a non-degenerate critical point of Λ_k and its Morse index equals the nodal surplus. In Corollary 4.6 it is shown that the Hessian of Λ_k is diagonal with respect to the decomposition (2). Therefore, the Morse index at $\varepsilon \in \{0, \pi\}^\beta$ is the number of segments in the 1-skeleton of the Boolean lattice that start at α and descend. By the monotonicity Lemma 5.1, this is the same as the number of segments whose endpoints have a lower eigenvalue, which is $|J_-(\alpha)|$. ■

7.2. A main tool that we will use in proving Proposition 3.5 is the local-global theorem of [11], which can be stated in a simplified manner as follows.

Theorem 7.3 ([11, Theorem 3.10]). *Suppose G is a simple, connected graph and $h \in \mathcal{S}(G)$ has a simple eigenvalue $\lambda_k(h)$ with a nowhere-vanishing eigenvector. Let $J \subset [\beta]$, let $\mathbb{T}_J \subset \mathbb{T}^\beta$ be the subtorus spanned by $\{e_j\}_{j \in J}$, and consider the restriction of Λ_k to the subtorus $\mathbb{T}_J$ (with $\Lambda_k(\alpha) = \lambda_k(\alpha * h)$ as before). Then, $\alpha = 0$ is a local minimum (resp. maximum) of Λ_k on $\mathbb{T}_J$ if and only if it is a global minimum (resp. maximum) on $\mathbb{T}_J$.*

The statements in [11] involve a different but equivalent graph model, and apply in a situation of greater generality, where the eigenvector is permitted to vanish at various vertices. We therefore provide the proof for Theorem 7.3, adapted to our situation, in Appendix A. Theorem 7.3 together with the monotonicity lemma gives the following.

Corollary 7.4. *Fix $h \in \mathcal{S}(G)$. Fix $\varepsilon \in \{0, \pi\}^\beta$ and write $h_\varepsilon = \varepsilon * h$. Let $\mathbb{T}_-(\varepsilon)$ denote the sub-torus of $\mathbb{T}^\beta$ that is spanned by those basis elements e_j for $j \in J_-(\varepsilon)$ and similarly for $\mathbb{T}_+(\varepsilon)$. Then,*

$$\begin{aligned} \lambda_k(\alpha * h_\varepsilon) &\leq \lambda_k(h_\varepsilon) \quad \text{for any } \alpha \in \mathbb{T}_-(\varepsilon) \\ \lambda_k(\alpha * h_\varepsilon) &\geq \lambda_k(h_\varepsilon) \quad \text{for any } \alpha \in \mathbb{T}_+(\varepsilon). \end{aligned}$$

7.5. Proof of Proposition 3.5

Suppose G is simple, connected, and has disjoint cycles, and suppose that h is generic in the sense of [GSC]. Let $\varepsilon, \varepsilon' \in \{0, \pi\}^\beta$. We need to show that

$$J_+(\varepsilon) = J_+(\varepsilon') \iff \varepsilon = \varepsilon'.$$

The definition of $J_{\pm}(\varepsilon)$ implicitly requires a choice of $k \in [n]$ and $h \in \mathcal{S}(G)$ so to be explicit we temporarily denote it $J_{\pm}(\varepsilon, k, h)$. For fixed ε, k , this set is constant (in h) on connected components of the open set $\mathcal{O}$ of $h \in \mathcal{S}(G)$ which satisfy condition [GSC], because the eigenvalues $\lambda_k(h), \lambda_k(\varepsilon * h)$ vary continuously with h . As a result, it is enough to prove the statement for $h \in \mathcal{B} \cap \mathcal{O}$ as this set is dense in $\mathcal{O}$ by Lemma 6.5. Recall that $h \in \mathcal{B} \cap \mathcal{O}$ if and only if it satisfies [GSC] and condition $(\spadesuit)$ which we repeat here:

$(\spadesuit)$ for each $k \in [n]$, the eigenvalue λ_k takes distinct values on distinct gauge-equivalence classes of signings of h .

Thus, we may assume that h satisfies [GSC] and $(\spadesuit)$. Given $\varepsilon, \varepsilon' \in \{0, \pi\}^{\beta}$, suppose $J_+(\varepsilon) = J_+(\varepsilon')$. Assume for the sake of contradiction that $\varepsilon \neq \varepsilon'$ so $\Lambda_k(\varepsilon) \neq \Lambda_k(\varepsilon')$. Assume that $\Lambda_k(\varepsilon) < \Lambda_k(\varepsilon')$ and let us show that there is ε'' such that $\Lambda_k(\varepsilon) > \Lambda_k(\varepsilon'') > \Lambda_k(\varepsilon')$ which provides the needed contradiction. Since $J_+(\varepsilon) = J_+(\varepsilon')$, then the intersection $\mathbb{T}_+(\varepsilon') \cap \mathbb{T}_-(\varepsilon)$ contains a signing, call it ε'' . Then, Corollary 7.4 implies $\Lambda_k(\varepsilon'') > \Lambda_k(\varepsilon')$ because $\varepsilon'' \in \mathbb{T}_+(\varepsilon')$. However, $\Lambda_k(\varepsilon'') < \Lambda_k(\varepsilon)$ because $\varepsilon'' \in \mathbb{T}_-(\varepsilon)$. ■

A. Proof of Theorem 7.3

A.1. We follow the proof in [11] but reorder the steps. Theorem 7.3 begins with a real symmetric matrix $h \in \mathcal{S}(G)$. Recall that the choice of edge $(r_j, s_j) \in \gamma_j$ determines a basis $e_1, e_2, \dots, e_{\beta}$ of $\mathbb{T}^{\beta} = \mathbb{R}^{\beta} / (2\pi i \mathbb{Z})^{\beta}$. The subset $J \subset [\beta]$ determines the subtorus $\mathbb{T}_J \subset \mathbb{T}^{\beta}$ which is spanned by the coordinates e_j for $j \in J$. We therefore have an analytic family of magnetic perturbations, $h_{\alpha} = \alpha * h$ for $\alpha \in \mathbb{T}_J$, and an eigenvalue function $\Lambda_k: \mathbb{T}_J \rightarrow \mathbb{R}$ defined by $\Lambda_k(\alpha) := \lambda_k(h_{\alpha})$. Since $\lambda = \lambda_k(h)$ is a simple eigenvalue, the function Λ_k is analytic near $\alpha = 0$ and is piecewise analytic on all of $\mathbb{T}_J$. We may choose the corresponding eigenvector ϕ of h to be real. By assumption, it is nowhere vanishing.

The point $\alpha = 0$ is a critical point of Λ_k . Assume it is a local minimum. Theorem 7.3 states that it is also a global minimum. (The case of a maximum can be proven analogously.) So we need to show

$$\lambda \leq \lambda_k(h_{\alpha}) \quad \text{for all } \alpha \in \mathbb{T}_J. \quad (11)$$

A.2. The proof in [11] involves several auxiliary matrices. Holding ϕ constant, the function $\langle \phi, h_{\alpha} \phi \rangle: \mathbb{T}_J \rightarrow \mathbb{R}$ has a critical point at $\alpha = 0$ (cf. equation (4)) and we set

$$\Omega = \frac{1}{2} \text{Hess}(\langle \phi, h_{\alpha} \phi \rangle) \Big|_{\alpha=0}.$$

The matrix Ω is a real diagonal $|J| \times |J|$ matrix. It is diagonal since each entry of h_α depends on at most one α_j coordinate, so $\frac{\partial^2 h_\alpha}{\partial \alpha_i \partial \alpha_j} = 0$ for $i \neq j$. It is real and invertible since its diagonal entries are

$$\Omega_{jj} = -h_{r_j s_j} \phi(r_j) \phi(s_j) \neq 0 \quad (12)$$

as calculated in (5). (Recall that both h and ϕ are real, and ϕ is nowhere-vanishing.)

For each $j \in J$, let $R_j(t)$ be the hermitian $n \times n$ matrix supported on the block

$$[r_j r_j, r_j s_j; s_j r_j, s_j s_j]$$

on which it is given by

$$R_j(t) = h_{r_j s_j} \begin{pmatrix} -\frac{\phi(s_j)}{\phi(r_j)} & e^{it} \\ e^{-it} & -\frac{\phi(r_j)}{\phi(s_j)} \end{pmatrix}.$$

To ease notation, let us assume that $J = \{1, 2, \dots, |J|\}$. Writing $\alpha = (\alpha_1, \alpha_2, \dots, \alpha_J) \in \mathbb{T}_J$, the sum

$$\sum_{j \in J} R_j(\alpha_j)$$

is a family of Hermitian $n \times n$ matrices depending on $\alpha \in \mathbb{T}_J$.

Define the real symmetric $n \times n$ (constant) matrix S by

$$S = h - \sum_{j \in J} R_j(0).$$

This collection of matrices satisfies the following properties:

(a) for any $\alpha = (\alpha_1, \alpha_2, \dots, \alpha_{|J|}) \in \mathbb{T}_J$,

$$h_\alpha = \alpha * h = S + \sum_{j \in J} R_j(\alpha_j);$$

(b) for any $j \neq j'$, $R_j(t)$ and $R_{j'}(t')$ commute for all t, t' ;

(c) $R_j(0)\phi = 0$ for every $j \in J$, and hence $S\phi = \lambda\phi$;

(d) $\det(R_j(t)) = 0$ so $R_j(t)$ has rank one;

(e) the semi-definite sign of $R_j(t)$ is independent of t since $\text{trace}(R_j(t)) = 2\Omega_{jj}$.

(See equation (12).) Let

$$m = |\{j \in J : -h_{r_i s_i} \phi(r_i) \phi(s_i) < 0\}| = \text{ind}(\Omega)$$

be the number of negative semi-definite R_j 's. Then the sum of these commuting rank-one matrices has m negative eigenvalues and $n - |J| > 0$ zero eigenvalues (recalling that $|J| \leq \beta < n$ by the assumption of disjoint cycles), so

$$\lambda_{m+1} \left(\sum_{j \in J} R_j(\alpha_j) \right) = 0 \quad \text{for all } \alpha \in \mathbb{T}_J.$$

The Weyl inequalities for $h_\alpha = S + \sum_{i=1}^{|J|} R_j(\alpha_j)$ may be expressed as follows:

$$\begin{aligned} \lambda_p(S) + \lambda_q(\Sigma R_j) &\leq \lambda_k(S + \Sigma R_j) \leq \lambda_s(S) + \lambda_r(\Sigma R_j) \\ (p + q \leq k + 1) &\hspace{10em} (k + n \leq r + s). \end{aligned}$$

Only the first inequality is required for the case of a local minimum. Taking $q = m + 1$ gives

$$\lambda_{k-m}(S) \leq \lambda_k(h_\alpha) \quad \text{for all } \alpha \in \mathbb{T}_J.$$

By (11), the proof of Theorem 7.3 now comes down to the following statement.

Lemma A.3. *If $\alpha = 0$ is a local minimum of $\Lambda_k(\alpha)$ then $\lambda_{k-m}(S) = \lambda_k(h) = \lambda$.*

The proof involves the next few paragraphs.

A.4. Holding ϕ constant gives a mapping $ih_\alpha\phi: \mathbb{T}_J \rightarrow \mathbb{C}^n$. Define B to be its derivative $B = iD(h_\alpha\phi)|_{\alpha=0}$. It is a real $n \times |J|$ matrix with

$$B_{vj} = \frac{\partial}{\partial \alpha_j} (h_\alpha \phi)_v \Big|_{\alpha=0} = \begin{cases} -h_{r_j s_j} \phi(r_j) & \text{if } v = s_j, \\ h_{s_j r_j} \phi(s_j) & \text{if } v = r_j, \\ 0 & \text{otherwise.} \end{cases}$$

A direct but messy calculation involving double subscripts as in [11, Lemma 2.7] gives

$$\sum_{j \in J} R_j(0) = B\Omega^{-1}B^T,$$

and, therefore,

$$S = h - B\Omega^{-1}B^T.$$

A.5. A primary insight in [11] is the identification of the generalized Schur complements in the real symmetric $(n + |J|) \times (n + |J|)$ matrix

$$M = \begin{pmatrix} h - \lambda & B \\ B^T & \Omega \end{pmatrix}. \tag{13}$$

These complements are defined to be

$$\begin{aligned} M/(h - \lambda) &= \Omega - B^T(h - \lambda)^+B, \\ M/\Omega &= (h - \lambda) - B\Omega^{-1}B^T = S - \lambda, \end{aligned}$$

where “+” denotes the Moore–Penrose pseudo-inverse.¹¹

¹¹The Moore–Penrose pseudo-inverse of a real symmetric matrix A is zero on $(\text{Im}(A))^\perp$ and is the inverse of the isomorphism $(\ker(A))^\perp \rightarrow \text{Im}(A)$.

Proposition A.6 ([11, Lemma 2.3]). *The Schur complement to $h - \lambda$ may be identified,*

$$M/(h - \lambda) = \frac{1}{2} \text{Hess}(\Lambda_k(0)).$$

Proof. The proof in [11] requires [21, Remark II.2.2, p. 81] but it is actually elementary and we provide it here for completeness. The Lemma is equivalent to the statement that

$$\frac{1}{2} \langle \eta, \text{Hess}(\Lambda_k(0))\eta \rangle = \langle \eta, \Omega\eta \rangle - \langle B\eta, (h - \lambda)^+ B\eta \rangle \quad \text{for all } \eta \in T_0 \mathbb{T}^J = \mathbb{R}^J.$$

To calculate $\langle \eta, \text{Hess}(\Lambda_k(0))\eta \rangle$, choose an analytic one parameter family α_t with $\eta = \dot{\alpha}(0)$ and write $h_t = \alpha_t * h$ with simple eigenvalue $\Lambda_k(\alpha_t)$ and normalized eigenvector ϕ_t (so that $\phi = \phi_0$). From (3) and (6), the second derivative is

$$\begin{aligned} \langle \eta, \text{Hess}(\Lambda_k(0))\eta \rangle &= \left. \frac{d^2}{dt^2} \Lambda_k(\alpha_t) \right|_{t=0} = \left. \frac{d^2}{dt^2} (\langle \phi_t, h_t \phi_t \rangle) \right|_{t=0} \\ &= \langle \phi, \ddot{h}\phi \rangle + 2\Re[\langle \phi, \dot{h}\dot{\phi} \rangle]_{t=0}, \end{aligned}$$

where

$$\dot{h} = \left. \frac{d}{dt} h_t \right|_{t=0}, \quad \dot{\phi} = \left. \frac{d}{dt} \phi_t \right|_{t=0},$$

and $\ddot{h} = \ddot{h}_t|_{t=0}$ (This is the formula from [21] that is referenced in [11].) The first term agrees with the first term in $2\langle \eta, (M/h - \lambda)\eta \rangle$:

$$\frac{1}{2} \langle \phi, \ddot{h}\phi \rangle = \frac{1}{2} \left. \frac{d^2}{dt^2} \langle \phi, h_t \phi \rangle \right|_{t=0} = \langle \eta, \Omega\eta \rangle.$$

The t -derivative of $ih_t\phi$ (keeping ϕ fixed) is

$$B\eta = i\partial_\eta(h_\alpha\phi) = i\dot{h}\phi$$

So, we need to compare

$$-\langle \eta, B^*(h - \lambda)^+ B\eta \rangle = -\langle \dot{h}\phi, (h - \lambda)^+ \dot{h}\phi \rangle$$

with $\langle \phi, \dot{h}\dot{\phi} \rangle = \langle \dot{h}\phi, \dot{\phi} \rangle$. From (3),

$$\dot{\phi} + (h - \lambda)^+ \dot{h}\phi = c\phi$$

for some constant c , because ϕ spans the (one-dimensional) kernel of $h - \lambda$. Taking the inner product with $\dot{h}\phi$ and using (4) with $\dot{\lambda} = 0$ gives

$$\langle \dot{h}\phi, \dot{\phi} \rangle + \langle \dot{h}\phi, (h - \lambda)^+ \dot{h}\phi \rangle = 0$$

as claimed. ■

A.7. Proof of Lemma A.3

The Hainsworth theorem for the matrix M in (13) gives,

$$\text{ind}(M) = \text{ind}(M/(h - \lambda)) + \text{ind}(h - \lambda) = \text{ind}(M/\Omega) + \text{ind}(\Omega),$$

which yields

$$\text{ind}(S - \lambda) = \text{ind}(M/(h - \lambda)) + k - 1 - m.$$

Since $\alpha = 0$ is a local minimum of Λ_k , Proposition A.6 gives

$$\text{ind}(M/(h - \lambda)) = 0.$$

Property (c) of the matrices $R_j(t)$ implies λ is an eigenvalue of S . Therefore,

$$\lambda_{k-m}(S - \lambda) = 0.$$

B. The BZ condition

The argument in Lemma 5.3 concerning eigenvalues with non-trivial multiplicity is essentially the same as that of [12, Theorem 1.5], which we state here for completeness because it is an important observation about singular critical points that may appear. We are interested in the Morse theory of the composition $\Lambda_k: \mathbb{T}^\beta \rightarrow \mathbb{R}$,

$$\mathbb{T}^\beta \rightarrow \mathcal{H}(G) \xrightarrow{\lambda_k} \mathbb{R}.$$

Fix $\alpha \in \mathbb{T}^\beta$ and suppose that $\lambda_k(\alpha * h)$ is an eigenvalue of multiplicity $m \leq \beta$. Let V denote the m -dimensional eigenspace. Consider the set of all Hermitian forms on V that are given by

$$\left\langle \phi, \frac{d}{dt}[(\alpha + tv) * h]v \right\rangle \quad \text{for } \phi, v \in V, \quad (14)$$

as v varies within $T_\alpha \mathbb{T}^\beta$. According to [12, Theorem 1.5], if there exists $v \in T_\alpha \mathbb{T}^\beta$ such that the form (14) is positive definite (which we refer to as the BZ condition), then the point $\alpha \in \mathbb{T}^\beta$ is topologically regular, meaning that for sufficiently small $\delta > 0$ the set $\mathbb{T}_{\leq \lambda - \delta}^\beta$ is a strong deformation retract of $\mathbb{T}_{\leq \lambda + \delta}^\beta$. (Here, $\mathbb{T}_{\leq t}^\beta = \{\alpha' \in \mathbb{T}^\beta : \Lambda_k(\alpha') \leq t\}$ and $\lambda = \Lambda_k(\alpha)$.)

Acknowledgments. The authors are grateful to R. Band and G. Berkolaiko for their encouragement. Part of the proof follows ideas developed by the first author together with Band and Berkolaiko in a joint unpublished project studying a similar question on quantum graphs. We are particularly grateful to G. Berkolaiko for his meticulous

reading of an earlier version of this manuscript, for pointing out an error in our original proof of Proposition 6.3, and for his many useful comments and suggestions. We would also like to thank Cynthia Vinzant and several anonymous referees for useful comments.

Funding. The first author was supported by the Simons Foundation Grant 601948, DJ.

Competing interests. The authors have no competing interests to declare that are relevant to the content of this article.

References

- [1] L. Alon, R. Band, and G. Berkolaiko, [Nodal statistics on quantum graphs](#). *Comm. Math. Phys.* **362** (2018), no. 3, 909–948 Zbl [1406.81037](#) MR [3845291](#)
- [2] L. Alon, R. Band, and G. Berkolaiko, [Universality of nodal count distribution in large metric graphs](#). *Exp. Math.* **33** (2024), no. 2, 301–335 Zbl [1542.05095](#) MR [4755238](#)
- [3] L. Alon, G. Berkolaiko, and M. Goresky, Smooth points of simple eigenvalues on the torus of magnetic perturbations of graphs. 2025, arXiv:[2505.17215v1](#)
- [4] L. Alon and M. Goresky, [Morse theory for discrete magnetic operators and nodal count distribution for graphs](#). *J. Spectr. Theory* **13** (2023), no. 4, 1225–1260 Zbl [1536.05284](#) MR [4707543](#)
- [5] L. Alon, D. Mikulincer, and J. Urschel, Nodal count for orthogonal ensembles. In preparation
- [6] L. Alon and J. Urschel, Average nodal count and the nodal count condition for graphs. 2024, arXiv:[2404.03151v1](#)
- [7] R. Band, [The nodal count \$\{0, 1, 2, 3, \dots\}\$ implies the graph is a tree](#). *Philos. Trans. R. Soc. Lond. Ser. A Math. Phys. Eng. Sci.* **372** (2014), no. 2007, article no. 20120504 Zbl [1352.34019](#) MR [3151080](#)
- [8] D. Beliaev and Z. Kereta, [On the Bogomolny-Schmit conjecture](#). *J. Phys. A* **46** (2013), no. 45, article no. 455003 Zbl [1278.82029](#) MR [3131565](#)
- [9] G. Berkolaiko, [A lower bound for nodal count on discrete and metric graphs](#). *Comm. Math. Phys.* **278** (2008), no. 3, 803–819 Zbl [1171.05356](#) MR [2373444](#)
- [10] G. Berkolaiko, [Nodal count of graph eigenfunctions via magnetic perturbation](#). *Anal. PDE* **6** (2013), no. 5, 1213–1233 MR [3125554](#)
- [11] G. Berkolaiko, Y. Canzani, G. Cox, and J. L. Marzuola, [A local test for global extrema in the dispersion relation of a periodic graph](#). *Pure Appl. Anal.* **4** (2022), no. 2, 257–286 Zbl [1505.35312](#) MR [4496087](#)
- [12] G. Berkolaiko and I. Zelenko, [Morse inequalities for ordered eigenvalues of generic self-adjoint families](#). *Invent. Math.* **238** (2024), no. 1, 283–330 Zbl [07921932](#) MR [4794596](#)
- [13] Y. Bilu and N. Linial, [Lifts, discrepancy and nearly optimal spectral gap](#). *Combinatorica* **26** (2006), no. 5, 495–519 Zbl [1121.05054](#) MR [2279667](#)

- [14] G. Blum, S. Gnutzmann, and U. Smilansky, [Nodal domains statistics: A criterion for quantum chaos](#). *Phys. Rev. Lett.* **88** (2002), article no. 114101
- [15] E. Bogomolny and C. Schmit, [Percolation model for nodal domains of chaotic wave functions](#). *Phys. Rev. Lett.* **88** (2002), article no. 114102
- [16] Y. Colin de Verdière, [Magnetic interpretation of the nodal defect on graphs](#). *Anal. PDE* **6** (2013), no. 5, 1235–1242 Zbl [1281.35090](#) MR [3125555](#)
- [17] R. Courant, Ein allgemeiner Satz zur Theorie der Eigenfunktionen selbstadjungierter Differentialausdrücke. *Gött. Nachr.* **1923**, 81–84 JFM [49.0342.01](#)
- [18] M. Fiedler, [Eigenvectors of acyclic matrices](#). *Czechoslovak Math. J.* **25(100)** (1975), no. 4, 607–618 Zbl [0325.15014](#) MR [0387308](#)
- [19] S. Gnutzmann, U. Smilansky, and J. Weber, [Nodal counting on quantum graphs](#). *Waves Random Media* **14** (2004), no. 1, S61–S73 Zbl [1063.81056](#) MR [2042545](#)
- [20] R. A. Horn and C. R. Johnson, [Matrix analysis](#). Second edn., Cambridge University Press, Cambridge, 2013 Zbl [1267.15001](#) MR [2978290](#)
- [21] T. Kato, [Perturbation theory for linear operators](#). Second edn., Grundlehren Math. Wiss. 132, Springer, Berlin etc., 1976 Zbl [0435.47001](#) MR [0407617](#)
- [22] P. Kuchment, [An overview of periodic elliptic operators](#). *Bull. Amer. Math. Soc. (N.S.)* **53** (2016), no. 3, 343–414 Zbl [1346.35170](#) MR [3501794](#)
- [23] A. W. Marcus, D. A. Spielman, and N. Srivastava, [Interlacing families I: Bipartite Ramanujan graphs of all degrees](#). *Ann. of Math. (2)* **182** (2015), no. 1, 307–325 Zbl [1316.05066](#) MR [3374962](#)
- [24] F. Rellich, [Perturbation theory of eigenvalue problems](#). Gordon and Breach Science Publishers, New York etc., 1969 Zbl [0181.42002](#) MR [0240668](#)
- [25] H. K. Wimmer, [Rellich’s perturbation theorem on Hermitian matrices of holomorphic functions](#). *J. Math. Anal. Appl.* **114** (1986), no. 1, 52–54 Zbl [0606.15004](#) MR [0829110](#)

Received 20 November 2024; revised 29 April 2025.

Lior Alon

Department of Mathematics, Massachusetts Institute of Technology,
Simons Building (Building 2), 77 Massachusetts Avenue, Cambridge, MA 02139, USA;
lioralon@mit.edu

Mark Goresky

School of Mathematics, Institute for Advanced Study, 1 Einstein Drive, Princeton,
NJ 08540-4907, USA; goresky@ias.edu