ANNALES

DE

L'INSTITUT FOURIER

Ralf RUECKRIEMEN

Recovering quantum graphs from their Bloch spectrum

Tome 63, n° 3 (2013), p. 1149-1176.

<http://aif.cedram.org/item?id=AIF_2013__63_3_1149_0>

Ann. Inst. Fourier, Grenoble
63, 3 (2013) 1149-1176

RECOVERING QUANTUM GRAPHS FROM THEIR BLOCH SPECTRUM

by Ralf RUECKRIEMEN (*)

Abstract. — We define the Bloch spectrum of a quantum graph to be the map that assigns to each element in the deRham cohomology the spectrum of an associated magnetic Schrödinger operator. We show that the Bloch spectrum determines the Albanese torus, the block structure and the planarity of the graph. It determines a geometric dual of a planar graph. This enables us to show that the Bloch spectrum indentifies and completely determines planar 3-connected quantum graphs.

Résumé. — Nous définissons le spectre de Bloch d'un graphe quantique comme la fonction qui assigne à chaque élément de la cohomologie de deRham le spectre d'un opérateur de Schrödinger magnétique associé. On montre que le spectre de Bloch détermine le tore d'Albanese, la structure de bloc et la planarité du graphe. Il détermine un dual géometrique d'un graphe planaire. Cela nous permet de montrer que le spectre de Bloch identifie et détermine complètement les graphes quantiques planaires 3-connexes.

1. Introduction

We consider finite combinatorial graphs in which each edge is equipped with a positive finite length, often called metric graphs. We define a Schrödinger operator on a metric graph. The pair of a metric graph together with a Schrödinger operator is called a quantum graph.

Quantum graphs are studied in mathematics and physics. They serve as simplified models in many settings involving wave propagation. The fact that they are essentially one-dimensional makes explicit computations

Keywords: quantum graphs, Schrödinger operators, spectrum, inverse spectral problem.
Math. classification: 35R30, 58J50, 58J53, 78A46, 81Q10, 58C40.
(*) I would like to thank my supervisor Carolyn Gordon for her extensive support. Without our frequent meetings and her numerous helpful and necessary suggestions this paper could not have been written. I'm also grateful to the reviewer, he made some excellent suggestions to improve the readability of this paper.

possible in various situations. On the other hand, the graph structure gives them enough complexity to be useful models. The papers [13], [15] and [3] provide an introduction and a survey of quantum graphs and the trace formulae that are often used to study them.

A function on a quantum graph consists of a function on each edge, where the edges are viewed as intervals. A Schrödinger operator acts on the space of all functions that are smooth on each edge and satisfy specified conditions at the vertices. We will impose the Kirchhoff vertex conditions, which require that the function be continuous and that the sum of the inward pointing derivatives on all edges incident at the vertex be zero. Kirchhoff conditions model a conservation of flow. A Schrödinger operator is a second order differential operator on each edge with leading term the standard Laplacian $-\left(\frac{\partial}{\partial x}\right)^2$. The first order part is called the magnetic potential.

The question to what degree the spectrum of the Laplace operator determines the underlying space was popularized my Kac in [10] in the manifold setting. The Schrödinger operator with zero magnetic potential is the standard Laplacian on a quantum graph. The relation between its spectrum or more generally the spectrum of a Schrödinger operator and the underlying graph is an active area of research. Various exact trace formulae relate the two, see for example [16],[12] or [3].

A quantum graph is determined by the spectrum of a single Schrödinger operator, if some genericity assumptions on the edge lengths in the quantum graph are made, [9]. This is not true without a genericity assumption on the edge lengths. Various examples of isospectral non-isomorphic quantum graphs exist, see for example [2], [9] or [1]. The last also proves a generalization of Sunada's theorem to construct isospectral quantum graphs.

The idea of this paper is to look at the spectra of an entire collection of Schrödinger operators, which we call the Bloch spectrum of a quantum graph.

The classical Bloch spectrum of a torus $\mathbb{R}^n/L$ assigns to each character $\chi : L \to \mathbb{C}^*$ the spectrum of the Euclidean Laplacian acting on the space of functions on $\mathbb{R}^n$ that satisfy $f(x+l) = \chi(l)f(x)$ for all $l \in L$. See, for example, [5] for inverse spectral results concerning the Bloch spectrum. As pointed out by Guillemin [8], the Bloch spectrum can also be interpreted as the collection of spectra of all operators $\nabla^*\nabla$ acting on sections of the trivial bundle, where ∇ is a connection with zero curvature. The set of these connections is given by $\nabla = (d + i\alpha)$ with α a harmonic 1-form on $\mathbb{R}^n/L$. (One may take α to be any closed 1-form, but the spectrum depends only on the cohomology class of α, so one may always assume α to be harmonic.) The correspondence with the classical notion is given by the association of the character $\chi(l) = e^{2\pi i\alpha(l)}$ to the harmonic form α, where now α is viewed as a linear functional on $\mathbb{R}^n$. This interpretation of the Bloch spectrum admits a generalization to arbitrary Hermitian line bundles over a torus, where now one considers all connections with, say, harmonic curvature, see [7].

Both interpretations of the Bloch spectrum can be carried over to quantum graphs. We use differential forms to define our operators, our approach is similar to the one in [15]. We will consider operators of the form $\Delta_\alpha = (d + 2\pi i\alpha)^*(d + 2\pi i\alpha)$ and vary the 1-form α. Similarly to the setting of flat tori the spectrum depends only on the equivalence class of α in $H^1_{dR}(G, \mathbb{R})/H^1_{dR}(G, \mathbb{Z})$. We can also define the Bloch spectrum using characters of the first fundamental group. We show that these two notions of the Bloch spectrum of a quantum graph are equivalent.

We want to see what information about the quantum graph can be retrieved from the Bloch spectrum without any genericity assumptions on the quantum graph. It is known (see section seven for details) that the spectrum of the standard Laplacian Δ_0 determines the dimension n of $H^1(G, \mathbb{R})$. Thus from that spectrum alone, we know that $H^1(G, \mathbb{R})/H^1(G, \mathbb{Z})$ is isomorphic as a torus (i.e., as a Lie group) to $\mathbb{R}^n/\mathbb{Z}^n$. Hence we can view the Bloch spectrum as a map that associates a spectrum to each $\alpha \in \mathbb{R}^n/\mathbb{Z}^n$. We ask the following question:

Suppose we are given a map that assigns a spectrum to each element α of $\mathbb{R}^n/\mathbb{Z}^n$ and we know these spectra form the Bloch spectrum of a finite quantum graph G. From this information, can one reconstruct G both combinatorially and metrically?

We will consider a generic α, i.e., one whose orbit is dense in the torus $\mathbb{R}^n/\mathbb{Z}^n$, and we will just consider the spectra associated to an interval in the orbit of $\alpha \in \mathbb{R}^n/\mathbb{Z}^n$.

Our main results are as follows.

Theorem 1.1. — *The Bloch spectrum determines the Albanese torus, $Alb(G) = H_1(G, \mathbb{R})/H_1(G, \mathbb{Z})$, of a quantum graph as a Riemannian manifold.*

The Albanese torus contains information about the length of cycles in the quantum graph and how they overlap. Note that the spectrum of a single Schrödinger operator does not determine the Albanese torus, there are examples of isospectral quantum graphs with different Albanese tori,

[2]. If the quantum graph is equilateral the Albanese torus also determines the complexity of the graph by a theorem in [11].

The block structure of a graph contains the broad structure of the graph, see definition 2.9 for the definition of *block structure*.

THEOREM 1.2. — *The Bloch spectrum determines the block structure of a quantum graph.*

The cycle space of a graph is closely related to its homology, we use some of its properties to show:

THEOREM 1.3. — *The Bloch spectrum determines whether or not a graph is planar.*

Planarity is not determined by the spectrum of a single Schrödinger operator, [2]. The information about the homology we read out from the Bloch spectrum allows us to construct a geometric dual of a planar quantum graph. We use it to show:

THEOREM 1.4. — *Planar* 3*-connected quantum graphs are completely determined by their Bloch spectrum.*

The plan of this paper is as follows. In the second section we collect various facts about combinatorial graphs that will be needed later on. In the third section we define differential forms on quantum graphs and use them to define the Schrödinger operators and the Bloch spectrum via differential forms. We define the Bloch spectrum via characters in section four and show that this definition is equivalent to the differential form version. We then define the Albanese torus of a quantum graph in the fifth section. We discuss a trace formula for quantum graphs, the key tool that allows us to get information about the quantum graph from the Bloch spectrum. In section seven we show that the Bloch spectrum determines the length of a shortest representative of each element in $H_1(G,\mathbb{Z})$, see theorem 7.7. This is the main theorem that relates the Bloch spectrum to the quantum graph. The other theorems are just consequences from this one. We then show that the Bloch spectrum determines the Albanese torus of a quantum graph. In section eight we use these properties to show that the Bloch spectrum determines the block structure and planarity of a quantum graph. If the graph is planar it determines a geometric dual of the graph. This information completely determines the underlying combinatorial graph from the Bloch spectrum if the graph is planar and 3-connected. In section nine we show that if we know the underlying combinatorial graph and it is planar and 3-connected then the Bloch spectrum determines the length of all edges in the graph, so we can recover the full quantum graph. In section ten we will treat disconnected graphs and show that our results still hold in this case.

2. Combinatorial graph theory

This chapter collects various basic facts about combinatorial graphs that will be required later. The material is mostly taken from [4], which provides an excellent introduction to the area.

All our graphs are finite and connected. We will treat the case of disconnected graphs in section ten. We allow loops and multiple edges. Let G be a graph. We will denote the set of vertices by V and the set of edges by $E(G)$ or E if there is no risk of confusion. Each edge has its two end vertices associated to it.

Remark 2.1. — We will assume throughout the paper that our graphs do not have vertices of degree 2. Once we pass to quantum graphs, two edges connected by a vertex of degree 2 with Kirchhoff boundary condition behave exactly the same way as a single longer edge does.

DEFINITION 2.2. — *A cycle in a graph is a closed walk that does not repeat any edges or vertices. Whenever we use the word cycle in this paper we mean it in this graph theoretical sense and not in a homological sense.*

DEFINITION 2.3. — *Let γ_1 and γ_2 be two oriented cycles in a graph. We say they have edges of positive overlap if they have an edge in common and pass through it in the same direction. We say they have edges of negative overlap if they have an edge in common and pass through it in opposite directions.*

Note that two cycles can have both edges of positive and negative overlap.

LEMMA 2.4. — *Every graph admits a basis of its homology that consists of cycles.*

Proof. — Pick a spanning tree of the graph. Associate to each edge of G not in the spanning tree the cycle that consists of this edge and the path in the spanning tree that connects its end points. This collection of cycles is a basis of the homology. □

DEFINITION 2.5. — *We call a graph with no leaves, that is vertices of degree 1, a leafless graph.*

2.1. The block structure of a graph

The block structure provides a broad view of the structure of a graph.

Definition 2.6. — *A graph G is called k-connected if any two vertices $v, v' \in V$ can be connected by k disjoint paths. The paths are called disjoint if they do not share any edges or vertices (apart from v and v').*

Definition 2.7. — *A vertex v in G is called a cut vertex if $G \setminus \{v\}$ is disconnected.*

Definition 2.8. — *Consider the set of all cycles in the graph. Declare two cycles equivalent if they share at least one edge. This generates an equivalence relation. We define the set of blocks to be the set of equivalence classes.* [(1)]

Note that all loops in the graph are blocks, all other blocks are 2-connected.

Definition 2.9. — *We define the block structure of a graph as follows. Each block in the graph is replaced by a small circle that we call a fat vertex. The cut vertices contained in this block correspond to the different attaching points on the fat vertex. For loops we interpret their vertex as the cut vertex where they are attached to the rest of the graph.*

All other blocks or remaining edges sharing one of the cut vertices with the original block are connected at the respective attaching point on the fat vertex.

It does not matter how the different attaching points are arranged around the fat vertex. We explicitly allow several fat vertices to be directly connected to each other without an edge in between. [(2)]

Example 2.10. — Figure 2.1 shows an example of a graph and its block structure. Note that for simplicity of recognition all blocks in the graphs are either loops or copies of the complete graph on 4 vertices, K_4.

Remark 2.11. — Any cycle in the graph is confined to a single block. Thus the vertices and edges in the block structure never form a cycle and the block structure has a tree like shape.

---

(1) This is a slight deviation from the standard definition. It is changed to allow graphs with loops and multiple edges. Edges that are not part of any cycle are not part of any block in our definition. Usually these edges are counted as blocks, too.

(2) Again this is a non standard definition. Our definition contains the same information about the graph as the standard one modulo the addition of loops and multiple edges.

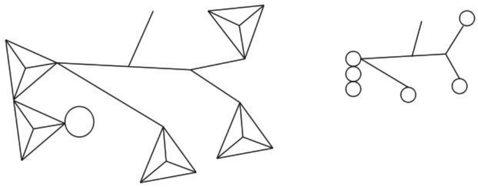

Figure 2.1. a quantum graph and its block structure

We will phrase the next two lemmata in the context of quantum graphs as we will need it later on.

Lemma 2.12. — *Let G be a quantum tree graph with no vertices of degree 2. Then the set of distances between any pair of leaves determines both the combinatorial tree graph underlying G and all individual edge lengths.*

Proof. — Given three leaves B_i, B_j and B_k the restriction of the tree to the paths between these leaves is shaped like a star. We will denote the length of the three branches by l_i, l_j and l_k. The distances between the leaves determine the quantities $l_i + l_j$, $l_i + l_k$ and $l_j + l_k$ and thus the three individual lengths l_i, l_j and l_k. This means that given a path between two leaves B_i and B_j and a third leaf B_k we can find both the point on the path from B_i to B_j where the paths from B_i and B_j to B_k branch away and the length of the path from this point to B_k.

We will use this fact repeatedly and proceed by induction on the number of leaves.

If there are only two leaves the tree consists of a single interval with length the distance between the two leaves.

Suppose we already have a quantum tree graph with leaves $B_1, \dots, B_{n-1}$. We now want to attach a new leaf B_n. We will first look at the leaves B_1 and B_2 and find the point on the path from B_1 to B_2 where the paths to B_n branch away. If this point is not a vertex of the tree, we create a new vertex and attach the leaf B_n on an edge of suitable length l_n. If this point is a vertex of the tree we know that the attachment point of B_n has to lie on the subtree branching away from the path from B_1 to B_2 starting at that vertex. Pick a leaf on this subtree, without loss of generality B_3, and look at the path from B_1 to B_3. We can again find the point on that path where the paths to B_n branch away. If this point is not a vertex of the tree we found the attachment point, otherwise we have reduced our search to an even smaller subtree. Continuing this process we will eventually end up with an attachment point on an edge or on a subtree that consists of a single vertex. □

LEMMA 2.13. — *Let G_0 be a 3-connected combinatorial graph and let G be a quantum graph with underlying combinatorial graph G_0. Then knowing G_0 and the length of each cycle determines the length of each edge in G.*

Proof. — Given an edge e there are at least 3 disjoint paths that connect its end vertices as G_0 is 3-connected. Thus there are two cycles in G_0 that share the edge e and its end vertices but otherwise are disjoint. Denote these two cycles by c_1 and c_2. Denote the closed walk $c_1 \setminus \{e\} \cup (-c_2 \setminus \{e\})$ by c_3. Since c_1 and c_2 are disjoint away from e the closed walk c_3 is a cycle. The length of e is given by $2L(e) = L(c_1) + L(c_2) - L(c_3)$ and thus is determined by the lengths of the cycles. □

2.2. Planarity of graphs

The edge space of a graph is the F_2-vector space over the set of (unoriented) edges of the graph. The cycle space is the subspace generated by cycles in the graph.

Given an embedding into $\mathbb{R}^2$ of a planar graph the faces of the embedding are the disconnected components of $\mathbb{R}^2 \setminus G$.

THEOREM 2.14. — *MacLane (1937)* [4]
A graph is planar if and only if its cycle space has a simple basis.

Simple means that each edge is part of at most 2 cycles in the basis.

COROLLARY 2.15. — *A graph is planar if and only if it admits a basis of its homology consisting of oriented cycles having no positive overlap.*

Proof. — Each cycle is confined to a single block of the graph and two cycles in different blocks share at most a single vertex and thus have zero overlap. Thus it is sufficient to prove the statement for 2-connected graphs.

Assume G is planar and 2-connected and choose an embedding into $\mathbb{R}^2$. The set of boundaries of faces with the exception of the outer face forms a basis of the homology that consists of cycles and is simple, see [4]. We orient all basis cycles counterclockwise. Then no two of them can run through the same edge in the same direction as no basis cycle can lie inside another basis cycle. Thus there are no edges of positive overlap.

Any basis of the homology where every basis element can be represented by a cycle in the graph gives rise to a basis of the cycle space consisting of exactly these cycles. Thus if the graph is not planar any basis of cycles of the homology is not simple by MacLanes theorem. Therefore there exists an edge that is part of three basis cycles. No matter how we orient these three cycles, two of them have to go through this edge with the same orientation and thus have edges of positive overlap. □

Definition 2.16. — *We call a basis without edges of positive overlap a non-positive basis of the graph and remark that a non-positive basis is always simple.*

If G is 2-connected and planar we can find a simple basis by picking the boundaries of faces. This proposition states that the converse is true, too.

Proposition 2.17. — [14] *Given a simple basis of the cycle space of a 2-connected planar graph there exists an embedding into $\mathbb{R}^2$ such that all basis elements are boundaries of faces.*

2.3. Dual graphs

Planar graphs have a notion of a dual graph. We will present two different ways of defining it and list some properties.

Definition 2.18. — *Given a planar graph G we associate to each embedding into the plane a geometric dual graph G^*. The vertices of G^* are the faces in the embedding of G. The number of edges joining to vertices in G^* is the number of edges that the corresponding faces in G have in common.*

Definition 2.19. — *A cut of a graph G is a subset of (open) edges S such that $G \setminus S$ is disconnected. A cut is minimal if no proper subset of S is a cut.*

Definition 2.20. — *Given a planar graph G, a graph G^* is an abstract dual of G if there is a bijective map $\psi : E(G) \to E(G^*)$ such that for any $S \subseteq E(G)$ the set S is a cycle in G if and only if $\psi(S)$ is a minimal cut in G^*.*

Proposition 2.21. — [4] *A planar graph can have multiple non isomorphic abstract duals. Any geometric dual of a planar graph is an abstract dual and vice versa. The dual of a planar graph is planar and G is an abstract dual of G^*. If G is 3-connected than G^* is unique up to isomorphism.*

Definition 2.22. — *We call two graphs G and H 2-isomorphic if there is a bijection between their edge sets that carries cycles to cycles. Note that this does not imply that the graphs are isomorphic.*

Example 2.23. — Figure 2.2 shows two graphs that are 2-isomorphic but not isomorphic. (In one of them the two vertices of degree 4 are adjacent, in the other one they are not.) The third graph is a common dual of them.

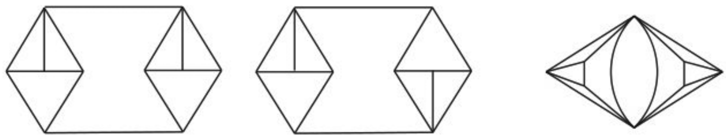

Figure 2.2. two 2-isomorphic graphs and one of their duals

Lemma 2.24. — *Two planar graphs G and H are 2-isomorphic if and only if they have the same set of abstract duals.*

Proof. — Let $\varphi : E(G) \to E(H)$ be a 2-isomorphism and let G^* be an abstract dual of G with edge bijection ψ. Then $\psi \circ \varphi^{-1}$ is an edge bijection that makes G^* an abstract dual of H.

Let G and H have the same abstract duals and let G^* be an abstract dual. Let ψ_1 be an edge bijection between G and G^* and let ψ_2 be an edge bijection between H and G^*. Then $\psi_2^{-1} \circ \psi_1$ is a 2-isomorphism between G and H. □

3. Differential forms

The concept of differential forms on a quantum graph was introduced in [6].

Let G be a quantum graph, let E and V be the set of edges and vertices. Let L_e denote the length of the edge e. Let $\{e \sim v\}$ denote the set of edges e adjacent to a vertex v.

Definition 3.1. — *A vector field X on G is a smooth vector field on each edge, seen as a closed interval. In particular a vector field is multivalued at the vertices.*

Let $\nu_{v,e}$ denote the outward unit normal for the edge e at the vertex v. Let X_1 be an auxiliary vector field that is real and has constant length 1 on all edges.

DEFINITION 3.2. — *A* 0*-form* f *on* G *is a function that is* C^∞ *on the edges, that is continuous and that satisfies the Kirchhoff boundary condition*

$$\sum_{e\sim v} \nu_{v,e}(f|_e) = 0$$

at all vertices $v \in V$. *We denote the space of* 0*-forms by* Λ^0.

DEFINITION 3.3. — *A* 1*-form* α *on* G *consists of a smooth* 1*-form* α_e *on each closed edge* e *such that* α *satisfies the boundary condition*

$$\sum_{e\sim v} \alpha_e(\nu_{v,e}) = 0$$

at all vertices $v \in V$. *We denote the space of* 1*-forms by* Λ^1.

DEFINITION 3.4. — *For a real* 1*-form* α *we define the operator* $d_\alpha : \Lambda^0 \to \Lambda^1$ *through the requirement*

$$(d_\alpha f)(X) := X(f) + 2\pi i \alpha(X) f$$

for all vector fields X. *We denote the operator* d_0 *by* d.

DEFINITION 3.5. — *We define a hermitian inner product on* Λ^0 *by*

$$(f,g) := \int_G f(x)\overline{g(x)}dx$$

DEFINITION 3.6. — *We define a hermitian inner product on* Λ^1 *by*

$$(\alpha,\beta) := \int_G \alpha(X_1)\overline{\beta(X_1)}dx = \sum_{e\in E}\int_0^{L_e} \alpha_e(X_1|_e)\overline{\beta_e(X_1|_e)}dx$$

This is clearly independent of the choice of the auxiliary vector field X_1.

We are now going to define the formal adjoint of d_α. Formally it should satisfy

$$(d_\alpha^* \beta, f) = (\beta, d_\alpha f)$$

for all $f \in \Lambda^0$. We have

$$\begin{aligned}(\beta, d_\alpha f) &= \int_G \beta(X_1)\overline{X_1(f)}dx - 2\pi i \int_G \beta(X_1)\overline{\alpha(X_1) f}dx \\ &= -\int_G X_1(\beta(X_1))\overline{f}dx + \sum_{v\in V}\overline{f(v)}\sum_{e\sim v}\beta_e(\nu_{v,e}) \\ &\qquad\qquad - 2\pi i \int_G \alpha(X_1)\beta(X_1)\overline{f}dx\end{aligned}$$

where we used integration by parts. The sum term vanishes because of the boundary condition on 1-forms. So we find that d_α^* satisfies

$$d_\alpha^*\beta = -X_1(\beta(X_1)) - 2\pi i\alpha(X_1)\beta(X_1) = d^*\beta - 2\pi i\alpha(X_1)\beta(X_1)$$

which again is independent of the choice of X_1.

Definition 3.7. — *For each edge $e \in E$ we define the Sobolov space $W_2(e)$ as the closure of $C^\infty([0, L_e])$ with respect to the norm $||f_e||_2^2 := \sum_{j=0}^{2} \int_0^{L_e} |f_e^{(j)}(x)|^2 dx$.*

We define the global Sobolov space $W_2(G)$ as the space of all functions f that are continuous on the entire graph and that satisfy $f|_e \in W_2(e)$ for all $e \in E$.

Definition 3.8. — *We define a Schrödinger type operator*

$$\Delta_\alpha := d_\alpha^* d_\alpha$$

on Λ^0. We extend its domain to

$$Dom(\Delta_\alpha) := \left\{ f \in W_2(G) \middle| \forall v \in V : \sum_{e\sim v} \nu_{v,e}(f|_e) = 0 \right\}$$

Theorem 3.9. — [6] *We have $H^1(G, \mathbb{C}) = \Lambda^1/d(\Lambda^0)$. Thus the definitions of 1-forms and 0-forms produce the expected deRham cohomology.*

Proposition 3.10. — *Let $\alpha \in \Lambda^1$ and $\psi \in \Lambda^0$ be real and let $\beta = \alpha + d\psi$. Let f be an eigenfunction of Δ_α with eigenvalue λ. Then $e^{-2\pi i\psi} f$ is an eigenfunction of Δ_β with the same eigenvalue. That is two 1-forms that differ by an exact 1-form have the same spectrum.*

Proof. — We have

$$\begin{aligned}
d_\beta^* d_\beta\left(e^{-2\pi i\psi} f\right) &= d_\beta^*\left(d(e^{-2\pi i\psi} f) + 2\pi i e^{-2\pi i\psi} f\alpha + 2\pi i e^{-2\pi i\psi} f d\psi\right) \\
&= d_\beta^*\left(e^{-2\pi i\psi} df + 2\pi i e^{-2\pi i\psi} f\alpha\right) \\
&= d_\beta^*\left(e^{-2\pi i\psi} d_\alpha f\right) \\
&= d^*\left(e^{-2\pi i\psi} d_\alpha f\right) - 2\pi i\alpha(X_1) e^{-2\pi i\psi} d_\alpha f(X_1) \\
&\qquad - 2\pi i d\psi(X_1) e^{-2\pi i\psi} d_\alpha f(X_1) \\
&= e^{-2\pi i\psi} d^* d_\alpha f - 2\pi i\alpha(X_1) e^{-2\pi i\psi} d_\alpha f(X_1) \\
&= e^{-2\pi i\psi} d_\alpha^* d_\alpha f
\end{aligned}$$

Thus f is an eigenfunction for Δ_α if and only if $e^{-2\pi i\psi} f$ is an eigenfunction for Δ_β with the same eigenvalue. □

Remark 3.11. — Note that $Spec_\alpha(G)$ depends only on the coset of $[\alpha]$ in $H^1_{dR}(G,\mathbb{R})/H^1_{dR}(G,\mathbb{Z})$.

Definition 3.12. — *Let $Spec_\alpha(G)$ be the spectrum of the operator Δ_α. We define the Bloch spectrum $Spec_{Bl}(G)$ of a quantum graph to be the map that associates to each $[\alpha]$ the spectrum $Spec_\alpha(G)$ where $[\alpha] \in H^1_{dR}(G,\mathbb{R})/H^1_{dR}(G,\mathbb{Z})$.*

Note that we assume that we only know $H^1_{dR}(G,\mathbb{R})/H^1_{dR}(G,\mathbb{Z})$ as an abstract torus without any Riemannian structure.

Definition 3.13. — *We say that two quantum graphs G and G' are Bloch isospectral if there is a Lie group isomorphism*

$$\Phi : H^1_{dR}(G,\mathbb{R})/H^1_{dR}(G,\mathbb{Z}) \to H^1_{dR}(G',\mathbb{R})/H^1_{dR}(G',\mathbb{Z})$$

such that

$$Spec_\alpha(G) = Spec_{\Phi(\alpha)}(G')$$

for all $[\alpha] \in H^1_{dR}(G,\mathbb{R})/H^1_{dR}(G,\mathbb{Z})$.

Note that if G is a tree graph the entire Bloch spectrum just consists of the spectrum of the Laplacian Δ_0 and thus does not contain any additional information.

4. The Bloch spectrum via characters

In this chapter we will introduce the Bloch spectrum using characters of the first fundamental group and show that the two notions are equivalent.

Let $\tilde{G}$ be the universal cover of G and let $\pi_1(G)$ denote the fundamental group. Then $\pi_1(G)$ acts by deck transformations on $\tilde{G}$. Let $\chi : \pi_1(G) \to \mathbb{C}^*$ be a character of $\pi_1(G)$.

We will study functions $\tilde{f} : \tilde{G} \to \mathbb{C}$ that are continuous and satisfy Kirchhoff boundary conditions at the vertices and that obey the transformation law

$$\tilde{f}(\gamma x) = \chi(\gamma)\tilde{f}(x)$$

for all $x \in \tilde{G}$ and $\gamma \in \pi_1(G)$. We refer to the space of these functions as $\Lambda^0_\chi(\tilde{G})$.

We associate to the character χ the spectrum of the standard Laplacian d^*d on $\tilde{G}$ restricted to functions in $\Lambda^0_\chi(\tilde{G})$, we will denote it by $Spec(G,\chi)$.

Definition 4.1. — *We call the map that associates to each character χ of $\pi_1(G)$ the spectrum $Spec(G,\chi)$ the π_1-spectrum of G.*

THEOREM 4.2. — *The Bloch spectrum $Spec_{Bl}(G)$ and the π_1-spectrum of a quantum graph are equal. There is a one-to-one correspondence $[\alpha] \mapsto \chi_\alpha$ between $H^1_{dR}(G,\mathbb{R})/H^1_{dR}(G,\mathbb{Z})$ and the set of characters of $\pi_1(G)$. It is given by*

$$\chi_\alpha(\gamma) = e^{-2\pi i \int_\gamma \alpha}$$

It induces the equality $Spec(G, \chi_\alpha) = Spec_\alpha(G)$.

Proof. — The integral does not depend on either the representative in $\pi_1(G)$ nor on the one in $H^1_{dR}(G,\mathbb{R})$ so this gives a well defined map. We also have $\chi_\alpha(\gamma_1 \cdot \gamma_2) = \chi_\alpha(\gamma_1)\chi_\alpha(\gamma_2)$ so this defines a character.

Let $f : G \to \mathbb{C}$ and let $\tilde{f} : \tilde{G} \to \mathbb{C}$ be the lift of f. Let $\tilde{\alpha}$ be the pullback of α. As $H^1(\tilde{G})$ is trivial $\tilde{\alpha}$ is exact and there exists a function $\tilde{\varphi} : \tilde{G} \to \mathbb{C}$ such that $\tilde{\alpha} = d\tilde{\varphi}$.

Let $\tilde{g}(x) := e^{-2\pi i \tilde{\varphi}(x)} \tilde{f}(x)$. We claim that $\tilde{g}$ is an eigenfunction in the π_1-spectrum if and only if $\Delta_\alpha f = \lambda f$. We need to show that $\tilde{g} \in \Lambda^0_{\chi_\alpha}(\tilde{G})$ and that $\Delta \tilde{g} = \lambda \tilde{g}$.

Let $\gamma \in \pi_1(G)$ and let $\tilde{\gamma}$ be the (unique) path in $\tilde{G}$ from x to γx. We have $\tilde{\varphi}(\gamma x) - \tilde{\varphi}(x) = \int_{\tilde{\gamma}} d\tilde{\varphi}$ by Stokes theorem. So we get

$$\tilde{g}(\gamma x) = e^{-2\pi i \tilde{\varphi}(\gamma x)} \tilde{f}(\gamma x) = e^{-2\pi i \int_{\tilde{\gamma}} \tilde{\alpha}} e^{-2\pi i \tilde{\varphi}(x)} \tilde{f}(x) = \chi_\alpha(\gamma) \tilde{g}(x)$$

By proposition 3.10 we have

$$\Delta \tilde{g} = \Delta e^{-2\pi i \tilde{\varphi}} \tilde{f} = e^{-2\pi i \tilde{\varphi}} \Delta_{\tilde{\alpha}} \tilde{f}$$

Thus $\tilde{g}$ is an eigenfunction with eigenvalue λ if and only if f is. □

Remark 4.3. — This theorem mirrors a similar result for tori, see [8].

5. The Albanese torus

DEFINITION 5.1. — *We call a 1-form α harmonic if $d^*\alpha \in \Lambda^0$ and $dd^*\alpha = 0$.*

LEMMA 5.2. — [6] *A 1-form α is harmonic if and only if $\alpha(X_1)$ is constant on all edges where X_1 is the auxiliary vector field of length one.*

LEMMA 5.3. — [6] *Any $\beta \in \Lambda^1$ admits a unique Hodge decomposition of the form*

$$\beta = d\psi + \tilde{\beta}$$

where $\psi \in \Lambda^0$ and $\tilde{\beta}$ is harmonic.

Thus each cohomology class has exactly one harmonic representative.

If β is real then so are ψ and $\tilde{\beta}$.

DEFINITION 5.4. — *We define an inner product on* $H^1(G,\mathbb{R})$ *by*

$$([\alpha],[\beta]) := (\tilde{\alpha},\tilde{\beta})$$

where $\tilde{\alpha}$ *and* $\tilde{\beta}$ *are the unique harmonic representatives of* $[\alpha]$ *and* $[\beta]$*.*

Kotani and Sunada defined the notion of the Albanese torus of a combinatorial graph in [11]. We will generalize this to quantum graphs. If the quantum graph is equilateral our definition recovers theirs.

Let $or(E)$ be the set of oriented edges, we call an element $b \in or(E)$ a bond. Let $\overline{b}$ denote a reversal of orientation. Let $o(b)$ and $t(b)$ be the origin and terminal vertex of a bond b.

Let A be an abelian group, the coefficients of the homology. Let $C_0(G,A)$ be the free A-module with generators in V. Let $C_1(G,A)$ be the A-module generated by $or(E)$ modulo the relation $\overline{b} = -b$. The boundary map $\partial : C_1(G,A) \to C_0(G,A)$ is defined by $\partial(b) := t(b) - o(b)$ and linearity. We then have $H_1(G,A) = ker(\partial)$.

We have the natural pairing $([\alpha],[p]) \mapsto \int_p \alpha$ for any $[\alpha] \in H^1(G,\mathbb{R})$ and $[p] \in H_1(G,\mathbb{R})$. This makes these two spaces dual to each other and induces an inner product on $H_1(G,\mathbb{R})$.

LEMMA 5.5. — *This inner product is equivalent to the one we get on* $H_1(G,\mathbb{R})$ *as a subspace of* $C_1(G,\mathbb{R})$ *with the inner product given by*

$$e \cdot e' = \begin{cases} l(e) & e = e' \\ -l(e) & e = \overline{e'} \\ 0 & otherwise \end{cases}$$

on edges and bilinear extension.

This might seem an awkward inner product if one thinks of vectors but the better analogy would be to think of characteristic functions of sets in $\mathbb{R}^n$ with an L^2 inner product.

Remark 5.6. — The inner product plays well with our notion of edges of positive and negative overlap in definition 2.3. The inner product of two cycles is equal to the difference between the length of the edges of positive and negative overlap.

DEFINITION 5.7. — *The Albanese torus of a quantum graph is the Riemannian torus given by*

$$Alb(G) = H_1(G,\mathbb{R})/H_1(G,\mathbb{Z})$$

with inner product as in 5.5. The Jacobian torus of a quantum graph is the Riemannian torus given by

$$Jac(G) = H^1(G,\mathbb{R})/H^1(G,\mathbb{Z})$$

with inner product as in 5.4. Note that these are dual tori.

6. A trace formula

The spectrum of a single Schrödinger type operator determines a trace formula. There are various different versions. All of them contain essentially the same information about the quantum graph. We will use the following.

THEOREM 6.1. — [12] *The spectrum $Spec_\alpha(G) = \{\lambda_n\}_n$ of the operator Δ_α determines the following exact wave trace formula.*

$$\sum_n \delta(\lambda - \sqrt{\lambda_n}) + \sum_n \delta(\lambda + \sqrt{\lambda_n})$$
$$= \frac{\mathcal{L}}{\pi} + \chi(G)\delta(\lambda) + \frac{1}{2\pi} \sum_{p \in PO} \left(\mathcal{A}_p(\alpha) e^{i\lambda l_p} + \overline{\mathcal{A}_p(\alpha)} e^{-i\lambda l_p} \right)$$

The statement of the trace formula in [12] is a lot more general than the one we give here, this special case with Kirchhoff boundary conditions can also be found in the survey paper [3].

The first sum is over all the eigenvalues including multiplicities, the δ are Dirac δ distributions.

$\mathcal{L}$ denotes the total edge length of the quantum graph. $\chi(G)$ denotes the Euler characteristic of the graph.

The second sum is over all periodic orbits, l_p denotes the length of a periodic orbit. A periodic orbit is an oriented closed walk in the quantum graph (without a fixed starting point).

The coefficients $\mathcal{A}_p(\alpha)$ are given by

$$\mathcal{A}_p(\alpha) = \tilde{l}_p e^{2\pi i \int_p \alpha} \prod_{b \in p} \sigma_{t(b)}$$

Here $\tilde{l}_p$ is the length of the primitive periodic orbit that p is a repetition of. The $e^{2\pi i \int_p \alpha}$ is the phase factor or 'magnetic flux'. The product is over the sequence of oriented edges or bonds in the periodic orbit. The vertex scattering coefficient $\sigma_{t(b)}$ at the terminal vertex $t(b)$ of each bond is given by $\sigma_{t(b)} = -\delta_{t(b)} + \frac{2}{\deg(t(b))}$. Here $\delta_{t(b)}$ is defined to be equal to one if the periodic orbit is backtracking at the vertex $t(b)$ and zero otherwise.

Remark 6.2. — The phase factor $e^{2\pi i \int_p \alpha}$ of a periodic orbit only depends on its homology class by Stokes theorem. For a contractible periodic orbit it is equal to 1.

COROLLARY 6.3. — [12] *The Fourier transform of this trace formula is given by:*

$$\begin{aligned} 1 + \sum_{n>0} e^{-il\sqrt{\lambda_n}} + \sum_{n>0} e^{+il\sqrt{\lambda_n}} \\ = 2\mathcal{L}\delta(l) + \chi(G) + \sum_{p \in PO} \mathcal{A}_p(\alpha)\delta(l - l_p) + \overline{\mathcal{A}_p(\alpha)}\delta(l + l_p) \end{aligned}$$

7. The homology of a quantum graph

In this chapter we will analyze the spectrum and the trace formula and extract information about the homology of the graph from it.

Before we state and prove the main theorem of this chapter we need a few definitions and a technical lemma.

DEFINITION 7.1. — *We call a periodic orbit minimal if it has minimal length within its homology class.*

Remark 7.2. — Note that in general a given element in the homology might have more than one minimal periodic orbit that represents it.

On the other hand, all closed walks that contain no edge repetitions, and in particular all cycles are minimal. Cycles are also the unique minimal periodic orbit in their homology class.

DEFINITION 7.3. — *We call a* 1*-form* α *generic if the image of the ray* $t\alpha$ *in the torus* $H^1_{dR}(G,\mathbb{R})/H^1_{dR}(G,\mathbb{Z})$ *is dense. The* α*'s with this property are dense. We pick and fix a single generic* α*.*

DEFINITION 7.4. — *To the fixed generic* α *we associate the following data.*

1. *Let* Ψ *be the linear map* $\Psi : H_1(G,\mathbb{Z}) \to \mathbb{R}$ *given by* $[p] \mapsto 2\pi \int_p \alpha$*. It associates to each periodic orbit its magnetic flux.*
2. *We call the absolute values of the magnetic fluxes* $\mu := |\Psi([p])| = 2\pi \left|\int_p \alpha\right|$ *the frequencies associated to* α*.*
3. *We will denote the length of the minimal periodic orbit(s) associated to a frequency* μ *by* $l(\mu)$*.*

Remark 7.5. — The map Ψ is two-to-one (except at zero) because we picked α to be generic. The set of all frequencies μ union their negatives $-\mu$ and zero forms a finitely generated free abelian subgroup of $\mathbb{R}$ that is isomorphic to $H_1(G,\mathbb{Z})$ via the map Ψ.

Lemma 7.6. — *Let f be a function that is a linear combination of several cosine waves with different (positive) frequencies.*

$$f(t)=\sum_{j=1}^{k}\nu_j\cos(\mu_j t)$$

Then the values $f(t)$ for $t\in[0,\varepsilon)$ determine both k and the individual frequencies $\mu_1,\ldots,\mu_k$.

Proof. — Assume without loss of generality that $0<\mu_1<\ldots<\mu_k$. We will show that we can determine μ_k and ν_k and then use induction. We will look at the collection of derivatives of f at $t=0$. We have

$$f^{(2n)}(0)=(-1)^n\sum_{j=1}^{k}\nu_j\mu_j^{2n}$$

There exists a unique number $\lambda>0$ such that

$$-\infty<\lim_{n\to\infty}\frac{f^{(2n)}(0)}{(-\lambda)^n}<\infty \qquad \textit{and} \qquad \lim_{n\to\infty}\frac{f^{(2n)}(0)}{(-\lambda)^n}\neq 0$$

and we have $\lambda=\mu_k^2$ and $\lim_{n\to\infty}\frac{f^{(2n)}(0)}{(-\lambda)^n}=\nu_k$. We can now look at the new function,

$$\tilde{f}(t):=f(t)-\nu_k\cos(\mu_k t)$$

repeat the process and determine μ_{k-1} and ν_{k-1}. After finitely many steps we will end up with the constant function 0. $\square$

The following theorem is the key link between the Bloch spectrum and the quantum graph. All other theorems are just consequences of this one.

Theorem 7.7. — *Given a generic α, see definition 7.3, the part of the Bloch spectrum $Spec_{t\alpha}(G)$ for $t\in[0,\varepsilon)$ determines the length of the minimal periodic orbit(s) of each element in $H_1(G,\mathbb{Z})$.*

Proof. — We will show we can read off the set of frequencies μ, see definition 7.4, associated to the generic α from the Bloch spectrum and determine the length $l(\mu)$ for each frequency.

We will look at the continuous family of 1-forms $\alpha(t)=t\alpha$ and the associated operators $\Delta_{\alpha(t)}$ for our fixed generic α. If we plug the eigenvalues of these operators into the Fourier transform of the trace formula we get a family of distributions. Each of these distributions is a locally finite sum of Dirac-δ-distributions (plus a constant term). The support of each of these δ-distributions is the length of the periodic orbit(s) it is associated to and thus depends only on the underlying quantum graph and not on the 1-form, see 6.3.

Any periodic orbit p that is homologically non-trivial has a corresponding partner which is the same closed walk with opposite orientation. Their coefficients are related by $\mathcal{A}_p(\alpha) = \overline{\mathcal{A}_{\overline{p}}(\alpha)}$ as the connectivity part and the length are the same and the magnetic flux changes sign. Thus for each such pair we would observe a factor of the form $2\mathrm{Re}\mathcal{A}_p(t\alpha) =: 2\mathrm{Re}\mathcal{A}_p(t)$ in the Fourier transform of the trace formulae for $\Delta_{\alpha(t)}$. We have

$$2\mathrm{Re}\mathcal{A}_p(t) = 2\mathrm{Re}\left(\tilde{l}_p e^{2\pi i t \int_p \alpha} \prod_{e \in p} \sigma_{t(e)}\right) = \nu \cos\left(2\pi \int_p \alpha t\right)$$

by theorem 6.1 where $\nu = 2\tilde{l}_p \prod_{e\in p} \sigma_{t(e)}$. So for each such pair of periodic orbits there is a magnetic flux $\Psi([p]) = 2\pi \int_p \alpha$ and a factor ν that is always non zero. Moreover the factor ν is positive if the periodic orbit contains no backtracks. As the magnetic flux appears in a cosine wave we can only know its absolute value, that is the frequency, see definition 7.4.

Pick a length of periodic orbits l. If we look at the family of 1-forms $\alpha(t)$ we get a continuous family of coefficients $\mathcal{A}^l(t) = \sum_{l_p = l} \mathcal{A}_p(t)$. As we did not make any assumptions on the underlying quantum graph there can be multiple periodic orbits with the same length. Thus each coefficient $\mathcal{A}^l(t)$ is a linear combination of a constant term and several cosine waves with different frequencies. The constant part comes from homologically trivial periodic orbits of length l. The cosine waves correspond to the homologically non-trivial periodic orbits of length l. We can now apply lemma 7.6 to the function $\mathcal{A}^l(t)$ and read off all the frequencies occurring at that length.

As we go through the different lengths in the spectra starting at zero we will pick up a collection of different frequencies. Each frequency will appear multiple times at different lengths since there are multiple periodic orbits that represent the same element in the homology and thus have the same frequency.

The frequency corresponding to a particular element in $H_1(G, \mathbb{Z})$ can only be realized by periodic orbits that represent this element in the homology, see remark 7.5. Going through the lengths starting at zero this frequency can appear at the earliest at the length of the corresponding minimal periodic orbit(s). The minimal periodic orbits need not be unique but as they are minimal they contain no backtracks. Thus their ν coefficients are all strictly bigger than 0 so their sum cannot vanish and the frequency will indeed appear in the coefficient $\mathcal{A}^l(t)$ at the minimal length. This gives us the length $l(\mu)$ associated to each frequency μ. □

Remark 7.8. — As we picked α to be generic the number of rationally independent frequencies is equal to $\dim H_1(G, \mathbb{Z})$. Thus we can observe from the number of rationally independent frequencies whether an arbitrary α is generic or not.

Remark 7.9. — Without any genericity assumptions on the edge lengths in the quantum graph it can happen that there are multiple non-minimal periodic orbits that are homologous and of the same length. We would not be able to distinguish them directly in the trace formula, it can even happen that their ν-coefficients cancel out and we would not observe them at all.

LEMMA 7.10. — *Given a frequency μ the following two statements are equivalent:*

(1) *The minimal periodic orbit associated to μ is a cycle in the graph.*
(2) *There are no two frequencies κ, κ', $\mu = |\kappa \pm \kappa'|$ with the property that $l(\kappa) + l(\kappa') \leqslant l(\mu)$.*

Proof. — We will prove both directions by contradiction.

Assume the minimal periodic orbit associated to μ is not a cycle, then it has to go through some vertex at least twice. Thus we can separate the periodic orbit into two shorter periodic orbits. Let κ and κ' be the frequencies associated to the two pieces. Then $\mu = |\kappa \pm \kappa'|$ and because the two pieces are not necessarily minimal $l(\kappa) + l(\kappa') \leqslant l(\mu)$.

Conversely, suppose μ admits a decomposition $\mu = |\kappa \pm \kappa'|$. Let c_μ, c_κ and $c_{\kappa'}$ denote the minimal periodic orbits associated to the frequencies. If we have $l(\kappa) + l(\kappa') = l(\mu)$ then $c_\mu = c_\kappa \cup c_{\kappa'}$ and c_κ and $c_{\kappa'}$ must have a vertex in common so c_μ is not a cycle. If we have $l(\kappa) + l(\kappa') < l(\mu)$ then the periodic orbits c_κ and $c_{\kappa'}$ are disjoint and c_μ realizes the connection between them so it uses the edges between them twice and is not a cycle. □

Remark 7.11. — Let μ_1 and μ_2 be two frequencies such that the associated minimal periodic orbits c_1 and c_2 are cycles. These cycles have an orientation induced from the 1-form α. The frequency $\mu_1 + \mu_2$ corresponds to the pair of periodic orbits that is homologous to $c_1 \cup c_2$ and $-(c_1 \cup c_2)$. Thus if $l(\mu_1+\mu_2) < l(\mu_1)+l(\mu_2)$ then c_1 and c_2 have edges of negative overlap. The frequency $|\mu_1 - \mu_2|$ corresponds to the pair of periodic orbits that is homologous to $c_1 \cup (-c_2)$ and $(-c_1) \cup c_2$. Thus if $l(\mu_1 - \mu_2) < l(\mu_1) + l(\mu_2)$ then c_1 and c_2 have edges of positive overlap.

Theorem 7.12. — *The Bloch spectrum of G determines the Albanese torus $Alb(G)$ as a Riemannian manifold.*

Proof. — Pick a minimal set of generators $\mu_1, \ldots, \mu_n$ of the group spanned by the frequencies such that the associated minimal periodic orbits are all cycles. Such a basis exists by lemma 2.4. Associate to them a set of vectors $v_1, \ldots, v_n$ satisfying $|v_i|^2 := l(\mu_i)$ and $2\langle v_i, v_j\rangle := l(\mu_i + \mu_j) - l(|\mu_i - \mu_j|)$ for all $i \neq j$. This uniquely determines a torus with spanning vectors $v_1, \ldots, v_n$.

If the cycles associated to μ_i and μ_j share no edges we have $l(\mu_i + \mu_j) = l(|\mu_i - \mu_j|) \geqslant l(\mu_i) + l(\mu_j)$ so the associated vectors are orthogonal.

If the cycles associated to μ_i and μ_j share edges the length $l(\mu_i + \mu_j)$ is twice the length of all edges of positive overlap plus the length of all edges that are part of one cycle but not the other. The length $l(|\mu_i - \mu_j|)$ is twice the length of all edges of negative overlap plus the length of all edges that are part of one cycle but not the other. Thus $l(\mu_i + \mu_j) - l(|\mu_i - \mu_j|)$ is twice the difference of the length of edges of positive overlap and the length of edges of negative overlap.

Therefore the torus is isomorphic to the Albanese torus of the quantum graph by lemma 5.5. □

The complexity of a graph is the number of spanning trees.

Corollary 7.13. — *If the quantum graph is equilateral the Bloch spectrum determines the complexity of the graph.*

Proof. — This follows directly from a theorem in [11]. For combinatorial graphs the complexity of the graph is given by $K(G) = \sqrt{vol(Alb(G))}$. The Albanese torus of an equilateral quantum graph is identical to the Albanese torus of the underlying combinatorial graph. □

Remark 7.14. — Leaves in a graph are invisible to the homology. So it is not clear whether the entire Bloch spectrum gives us any more information about them than the spectrum of a single Schrödinger type operator. There are examples of tree graphs that are isospectral for the standard Laplacian, see for example [9].

8. Determining graph properties

We will now use the information gained in the last section and translate it into graph properties that are determined by the Bloch spectrum.

Remark 8.1. — The Albanese torus distinguishes the isospectral examples of van Below in [2]. Thus the spectrum of a single Schrödinger type operator does not determine the Albanese torus. In one of the two quantum graphs two periodic orbits of length 3 can be composed to get a periodic orbit of length 4. Thus the lattice that corresponds to the Albanese torus contains two vectors of length 3 whose sum has length 4. In the other graph this is not the case. In particular these two quantum graphs are not Bloch isospectral by 7.12.

8.1. The block structure

THEOREM 8.2. — *One can recognize the block structure, see definition 2.9, of a leafless graph from the Bloch spectrum. It also determines the dimension of the homology of each block.*

Proof. — Pick a minimal set of generators $\mu_1, \dots, \mu_n$ of the group spanned by the frequencies such that the associated minimal periodic orbits are all cycles. A cycle is necessarily contained within a single block, see 2.11. Declare two generators equivalent if the associated cycles share edges regardless of orientation. This generates an equivalence relation. Let $\mathfrak{B}$ be the set of equivalence classes, it corresponds to the set of blocks of G, see 2.9. The number of generators in each equivalence class is the dimension of the homology of that block.

Let $B_1, B_2 \in \mathfrak{B}$. Let $\{\mu_i^j\}_i$ be the subset of frequencies that is B_j, $j = 1, 2$. Then we can find the distance between the two blocks by computing

$$d(B_1, B_2) := \frac{1}{2} \min_{i,i'} \left(l(\mu_i^1 + \mu_{i'}^2) - l(\mu_i^1) - l(\mu_{i'}^2) \right)$$

That is we compute the distance between any basis cycle in one block to any basis cycle in the other and minimize over all pairs of basis cycles in the blocks. This distance is zero if and only if the blocks share a vertex.

We will now set up a situation where we can apply lemma 2.12. To do so we need to find out which blocks are leaves in the block structure and which ones are inner vertices. We will then cut the block structure into smaller pieces such that all blocks are leaves in the smaller pieces.

Whenever we have a triple of blocks satisfying $d(B_2, B_3) > d(B_1, B_2) + d(B_1, B_3)$, that is a failure of the triangle inequality, we know that B_1 has to be an inner vertex in the block structure of G. The path between the blocks B_2 and B_3 has to pass through B_1 and use some edges within the block B_1. Once we have identified a block, say B_1, as an inner block we can separate the remaining blocks into groups depending on where the path from the block to B_1 is attached on B_1. If $d(B_i,B_j) > d(B_1,B_i)+d(B_1,B_j)$ then the paths from B_1 to B_i and B_j are attached at different cut vertices of B_1, if $d(B_i,B_j) \leqslant d(B_1,B_i) + d(B_1,B_j)$ they are attached at the same cut vertex. Within each of these groups the block B_1 is a leaf in the block structure.

Thus we have cut the initial block structure into several smaller pieces each of them including B_1 and B_1 is a leaf in each of them. We can repeat this process of identifying an inner block and cutting the block structure into smaller pieces on each of these pieces until all the pieces have no inner block vertices. This reduces the problem to recovering the block structure of graphs where all blocks are leaves.

All remaining inner vertices have to be vertices of the initial graph G and thus have degree at least 3. As G is leafless all leaves in the block structure are fat vertices. Hence we can recover the block structure of each of the smaller pieces by using lemma 2.12. We can then find the block structure of the entire graph by gluing the pieces together at the inner blocks. □

8.2. Planarity and dual graphs

THEOREM 8.3. — *The Bloch spectrum determines whether or not a graph is planar.*

Proof. — The homology admits infinitely many bases, but a graph has only finitely many cycles and thus there are only finitely many bases consisting of cycles. Thus there are only finitely many minimal sets of generators $\mu_1,\dots,\mu_n$ of the group spanned by the frequencies such that the minimal periodic orbits associated to them are all cycles. Given such a basis we can choose for each generator to either keep the orientation induced by α or choose the reverse orientation. Denote the basis elements with a choice of orientation by $\gamma_1 = \pm\mu_1,\dots,\gamma_n = \pm\mu_n$. For any pair of oriented basis elements γ_i and γ_j we can check whether the associated cycles have edges of positive overlap by checking whether $l(|\gamma_i+\gamma_j|)-l(|\gamma_i|)-l(|\gamma_j|) > 0$, see 7.11. Thus we can check whether the $\gamma_1,\dots,\gamma_n$ correspond to a basis of the homology that consists of oriented cycles having no positive overlap. The graph is planar if and only if we can find such a basis by theorem 2.15. □

Remark 8.4. — Planarity is a property that is not determined by the spectrum of a single Schrödinger type operator. There is an example of two isospectral quantum graphs in [2] where one is planar and the other one is not.

If the graph is planar we will fix a non-positive basis $\gamma_1, \ldots, \gamma_n$ (see definition 2.16) coming from the frequencies $\mu_1, \ldots, \mu_n$. We know that the basis elements are the boundaries of the inner faces in a suitable embedding of the graph by lemma 2.17. We will use this fact to construct an abstract dual of the graph.

Theorem 8.5. — *The Bloch spectrum of a planar, 2-connected graph determines a dual of the graph. Thus the Bloch spectrum determines planar, 2-connected graphs up to 2-isomorphism (see lemma 2.24).*

Before we show this we need two lemmata.

Lemma 8.6. — *Let G be planar and 2-connected. In the embedding where the $\gamma_1, \ldots, \gamma_n$ are the boundaries of the inner faces the boundary of the outer face is given by*

$$\gamma_0 := -\sum_{l=1}^{n} \gamma_l$$

The sign orients it so that it does not have edges of positive overlap with any of the γ_l.

Lemma 8.7. — *Let G be planar and 2-connected. Then we can determine the number of edges that any two of the cycles $\gamma_0, \ldots, \gamma_n$ have in common.*

Proof. — Recall that $\mu_i = |\gamma_i|$. If $l(\mu_i + \mu_j) \geqslant l(\mu_i) + l(\mu_j)$ then γ_i and γ_j share no edges. We will assume from now on that $l(\mu_i + \mu_j) < l(\mu_i) + l(\mu_j)$. Suppose γ_i and γ_j share k edges or single vertices.

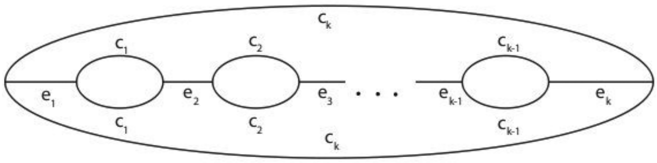

Figure 8.1. the graph with the two cycles γ_i and γ_j

Figure 8.1 shows the graph G. Here γ_i and γ_j bound the two big faces and the e_l are the edges or single vertices these two cycles share. The remainder of the graph is contained inside the small cycles labeled c_l, $l = 1, \ldots, k-1$ and outside the big cycle c_k.

We can decompose the minimal periodic orbit associated to the frequency $\mu_i + \mu_j$ into several cycles $c_1, \ldots, c_k$ by applying lemma 7.10 repeatedly.

These cycles do not share any edges, they can share vertices. If the decomposition yields k cycles, then γ_i and γ_j have k distinct components in common. Each of these components is either a single edge or a vertex. Whenever it is a vertex that means that two of the c_l have this vertex in common and thus have distance zero from each other. As we can check for any pair c_l and $c_{l'}$ whether $l(c_l + c_{l'}) = l(c_l) + l(c_{l'})$ we can find all instances where this happens. All remaining common components then must correspond to a common edge of γ_i and γ_j. □

Remark 8.8. — Lemma 8.7 is false without the planarity assumption. There exist two cycles γ_1, γ_2 in $K_{3,3}$ that share 3 edges but $\gamma_1 \cup \gamma_2$ is homologous to a single cycle. These two cycles have no edges of positive overlap.

Proof. — of theorem 8.5
The cycles $\gamma_0, \dots, \gamma_n$ are the set of all boundaries of faces in a suitable embedding of the graph. Therefore they are the vertices of a geometric dual. By lemma 8.7 we know the number of edges any two of these faces have in common, which corresponds to the number of edges between the two vertices in the geometric dual. □

The particular geometric dual we get from this process depends on the non-positive basis we have chosen.

Corollary 8.9. — *The Bloch spectrum identifies and determines planar, 3-connected graphs combinatorially.*

Proof. — Note that 3-connected implies 2-connected, see definition 2.6. A graph is 2-connected if its block structure consists of a single fat vertex. We have shown that we can identify 2-connected planar graphs in theorem 8.3. We found a geometric dual of a 2-connected planar graph in theorem 8.5. If the dual of the dual is 3-connected it will be unique and therefore isomorphic to the original graph. □

9. Determining the edge lengths

In this chapter we will show that we can recover all the edge lengths of a 3-connected planar graph if we know the underlying combinatorial graph.

The Bloch spectrum only gives us a map from the abstract torus $H^1(G,\mathbb{R})/H^1(G,\mathbb{Z})$ to the spectra. By theorem 7.7 we know the length of the minimal periodic orbit(s) associated to each element in $H_1(G,\mathbb{Z})$.

Here we need a little more, we want to associate the lengths we get from the Bloch spectrum with the periodic orbits in the combinatorial graph.

When we construct a dual graph in theorem 8.5 we can keep track of the lengths associated to the minimal set of generators $\mu_1, \dots, \mu_n$ of the group spanned by the frequencies. These frequencies correspond to the vertices of the dual graph. The vertices of the dual graph then correspond to a set of cycles in the dual of the dual that generates the homology. If the graph is 3-connected the dual of the dual is isomorphic to the original graph so we can associate the frequencies and their lengths to the periodic orbits in the graph.

Theorem 9.1. — *The Bloch spectrum identifies and completely determines 3-connected planar quantum graphs.*

Proof. — We have already shown in corollary 8.9 that the Bloch spectrum identifies 3-connected planar quantum graphs and determines their underlying combinatorial graph. All that remains to be shown is that we can determine all the edge lengths. By the remarks above the Bloch spectrum determines the length of all cycles in the quantum graph. We can now apply lemma 2.13 to determine the length of all the edges. ☐

Remark 9.2. — One can show that given the Bloch spectrum and the underlying combinatorial graph of an arbitrary quantum graph one can associate the frequencies from the Bloch spectrum and their lengths to the closed walks in the combinatorial graph. However, using this information to determine all the individual edge lengths is more delicate.

10. Disconnected graphs

If we do not assume that the quantum graph is connected we get a component-wise version of theorem 7.7.

Proposition 10.1. — *Let G be a quantum graph that may or may not be connected. Then the spectrum of the standard Laplacian Δ_0 determines the number of connected components. Denote the connected components by $G_1, \dots, G_k$.*

Given a generic α, see definition 7.3, the part of the Bloch spectrum $Spec_{t\alpha}(G)$ for $t \in [0, \varepsilon)$ determines the groups $H_1(G_i, \mathbb{Z})$ and the length of the minimal periodic orbit(s) of each element in $H_1(G_i, \mathbb{Z})$ for each component $i = 1, \dots, k$.

Proof. — The multiplicity of the eigenvalue 0 of the standard Laplacian Δ_0 is equal to the number of connected components.

The trace formula in theorem 6.1 still holds for disconnected graphs. The two sums over eigenvalues and periodic orbits are just unions over the connected components. The total edge length of the quantum graph is additive and the Euler characteristic is well defined for disconnected graphs, too.

Thus we can copy most of the proof of theorem 7.7 verbatim and read out a set of frequencies from the Bloch spectrum. Every frequency we get is associated to a single periodic orbit that belongs to only one of the connected components. If the sum of two frequencies is a frequency, then these two frequencies belong to the same connected component of G. If it is not they belong to different connected components. Thus the set of frequencies (union their negatives and zero) will not form one finitely generated free abelian subgroup of $\mathbb{R}$ that is isomorphic to $H_1(G,\mathbb{Z})$. Instead it will form k disjoint (apart from zero) finitely generated free abelian subgroups of $\mathbb{R}$ that are isomorphic to the $H_1(G_i,\mathbb{Z})$ for $i=1,\ldots,k$.

We can now assign a length to each frequency the same way as in theorem 7.7. □

As all our subsequent theorems are just consequences of theorem 7.7 they also hold component-wise.

Corollary 10.2. — *The theorems 7.12, 8.2, 8.3, 8.5 and 9.1 all hold component-wise.*

BIBLIOGRAPHY

[1] R. Band, O. Parzanchevski & G. Ben-Shach, "The isospectral fruits of representation theory: quantum graphs and drums", *J. Phys. A* **42** (2009), no. 17, p. 175202, 42.

[2] J. von Below, "Can one hear the shape of a network?", in *Partial differential equations on multistructures (Luminy, 1999)*, Lecture Notes in Pure and Appl. Math., vol. 219, Dekker, New York, 2001, p. 19-36.

[3] J. Bolte & S. Endres, "Trace formulae for quantum graphs", in *Analysis on graphs and its applications*, Proc. Sympos. Pure Math., vol. 77, Amer. Math. Soc., Providence, RI, 2008, p. 247-259.

[4] R. Diestel, *Graph theory*, third ed., Graduate Texts in Mathematics, vol. 173, Springer-Verlag, Berlin, 2005, xvi+411 pages.

[5] G. Eskin, J. Ralston & E. Trubowitz, "On isospectral periodic potentials in $\mathbf{R}^n$. II", *Comm. Pure Appl. Math.* **37** (1984), no. 6, p. 715-753.

[6] B. Gaveau & M. Okada, "Differential forms and heat diffusion on one-dimensional singular varieties", *Bull. Sci. Math.* **115** (1991), no. 1, p. 61-79.

[7] C. S. Gordon, P. Guerini, T. Kappeler & D. L. Webb, "Inverse spectral results on even dimensional tori", *Ann. Inst. Fourier (Grenoble)* **58** (2008), no. 7, p. 2445-2501.

[8] V. Guillemin, "Inverse spectral results on two-dimensional tori", *J. Amer. Math. Soc.* **3** (1990), no. 2, p. 375-387.

[9] B. Gutkin & U. Smilansky, "Can one hear the shape of a graph?", *J. Phys. A* **34** (2001), no. 31, p. 6061-6068.

[10] M. Kac, "Can one hear the shape of a drum?", *Amer. Math. Monthly* **73** (1966), no. 4, part II, p. 1-23.

[11] M. Kotani & T. Sunada, "Jacobian tori associated with a finite graph and its abelian covering graphs", *Adv. in Appl. Math.* **24** (2000), no. 2, p. 89-110.

[12] T. Kottos & U. Smilansky, "Periodic orbit theory and spectral statistics for quantum graphs", *Ann. Physics* **274** (1999), no. 1, p. 76-124.

[13] P. Kuchment, "Quantum graphs: an introduction and a brief survey", in *Analysis on graphs and its applications*, Proc. Sympos. Pure Math., vol. 77, Amer. Math. Soc., Providence, RI, 2008, p. 291-312.

[14] B. Mohar & C. Thomassen, *Graphs on surfaces*, Johns Hopkins Studies in the Mathematical Sciences, Johns Hopkins University Press, Baltimore, MD, 2001, xii+291 pages.

[15] O. Post, "First order approach and index theorems for discrete and metric graphs", *Ann. Henri Poincaré* **10** (2009), no. 5, p. 823-866.

[16] J.-P. Roth, "Le spectre du laplacien sur un graphe", in *Théorie du potentiel (Orsay, 1983)*, Lecture Notes in Math., vol. 1096, Springer, Berlin, 1984, p. 521-539.

Manuscrit reçu le 10 février 2011,
accepté le 29 mars 2011.

Ralf RUECKRIEMEN
Dartmouth College,
Kemeny Hall,
Hanover, 03755 NH, USA
ralf@rueckriemen.de