Laplacians and spectrum for singular foliations [1]

by Iakovos Androulidakis

National and Kapodistrian University of Athens
Department of Mathematics
Panepistimiopolis
GR-15784 Athens, Greece

e-mail: `iandroul@math.uoa.gr`

Abstract

We survey Connes' results on the longitudinal Laplace operator along a (regular) foliation and its spectrum, and discuss their generalization to any singular foliation on a compact manifold. Namely, we prove that the Laplacian of a singular foliation is an essentially self-adjoint operator (unbounded) and has the same spectrum in every (faithful) representation, in particular in L^2 of the manifold and L^2 of a leaf. We also discuss briefly the relation of the Baum-Connes assembly map with the calculation of the spectrum.

Introduction

Connes' Noncommutative Geometry programme has been very succcessful in the study of the leaf space of a foliation, whose topology is usually quite pathological. In fact, the foliation C^*-algebra plays the role of the space of continuous functions on the leaf space. Connes developed a longitudinal pseudodifferential calculus which allows the definition of a leafwise Laplacian for any (regular) foliation. Moreover, using the K-theory of $C^*(\mathcal{F})$ he managed to calculate its spectrum in some cases. Recall that in the case of a (compact) manifold, the spectrum of the Laplacian carries topological information; whence Connes' results are the first steps towards extracting similar information for much more pathological spaces.

In this note we overview these results and discuss their generalization to the singular case, namely to *any* singular foliation on a compact manifold M. We show that for any singular foliation $(M, \mathcal{F})$ the Laplacian defined in [2] as an unbounded multiplier of $C^*(\mathcal{F})$ is essentially self-adjoint in every (faithful) representation and has the same spectrum. In particular for its representations on $L^2(M)$ and L^2 of a leaf. The main ingredient that makes this possible is the holonomy groupoid of the (singular) foliation constructed in [1] and developed in [4], [5].

Regarding the calculation of the spectrum, notice that Connes' method requires the understanding of the "shape" of $K_0(C^*(\mathcal{F}))$, and the Baum-Connes conjecture provides a prediction for this. At the end of this sequel we briefly discuss a conjecture about the Baum-Connes assembly map associated with a singular foliation.

---

[1]AMS subject classification: 53C12, 35J05 47A10, 53C29, 22A22 . Keywords: Laplacian, singular foliation, holonomy.

1 Laplacians on Regular Foliations

We begin with an account of (regular) foliations (§1.1) and the associated holonomy groupoid (§1.2). The objective of this note (Laplacians and their spectrum) is introduced in §1.3 with the case of the Kronecker foliation.

1.1 Regular Foliations

A *foliation* is usually understood as a partition of a smooth manifold to connected and immersed maximal submanifolds (leaves) of *equal* (constant) dimension. In this sequel we are adding the adjective "regular" to such foliations, in order to distinguish them from the"singular" ones, namely those whose leaves may have non-constant dimension[2]. One archetype of a regular foliation to bear in mind is $\mathbb{R}^2$ foliated by parallel lines of slope $\theta \in \mathbb{R}$. This naturally provides the manifold with a sense of *"longitudinal"* and *"transversal"* direction. More strictly, we have the following characterization of regular foliations:

Characterization of regular foliations 1.1. *A regular foliation $\mathcal{F}$ on a smooth manifold M is characterized by an open cover of M by foliation charts of the form $\Omega = U \times T$, where $U \subseteq \mathbb{R}^p$ and $T \subseteq \mathbb{R}^q$. T is the transverse and U is the longitudinal direction. The transition functions are of the form*

$$f(u,t) = (g(u,t), h(t))$$

namely their second variable depends only on the transversal direction.

Here p is the dimension of the leaf and q is the codimension of the foliation (the dimension of M is $p + q$).

However, in general foliations appear in the study of dynamical systems, so in this context the following equivalent characterization of regular foliations on a compact manifold is more useful:

Characterization of regular foliations 1.2. *A regular foliation on a compact manifold M is the unique projective $C^\infty(M)$-submodule $\mathcal{F}$ of vector fields of M which are tangent to the leaves. This sub-module is also closed by the Lie bracket, namely if $X, Y \in \mathcal{F}$ then $[X, Y] \in \mathcal{F}$.*

Note that the Frobenius theorem ensures that a module $\mathcal{F}$ such as in 1.2 integrates to a partition to leaves. Moreover, the Serre-Swan theorem identifies $\mathcal{F}$ with the module of sections of a vector bundle F over M; namely the bundle of vectors in TM which are tangent to the leaves. We stress the uniqueness of the bundle F: It is completely determined by the partition to leaves. Let us give some examples:

Examples 1.3. Let M be a smooth manifold.

a) The orbits of the action of a Lie group G. The sub-module $\mathcal{F}$ is the image of the infinitesimal action $\mathfrak{g} \to \mathcal{X}(M)$ (where $\mathfrak{g}$ is the Lie algebra of G). Of course, not all Lie group actions give rise to regular foliations (think of the action of $GL(2)$ on $\mathbb{R}^2$).

---

[2]In fact, if "foliations" stands for the union of the regular with the singular case, then singular foliations appear much more often in mathematics than the regular ones appear in the literature.

b) If M is a Poisson manifold then it is naturally foliated; $\mathcal{F}$ is generated by the Hamiltonian vector fields and the leaves are symplectic manifolds. It turns out that the Poisson structure is completely determined by the symplectic foliation. Again, this foliation is not always regular (*e.g.* consider the linear Poisson structure on the Lie algebra dual of $SU(2)$; the symplectic leaves are the orbits of the coadjoint action; Identifying the Lie algebra dual with $\mathbb{R}^3$ these orbits are concentric spheres around the origin, so they present a singularity at the origin.)

The meaning of both the above examples was to exhibit that singular foliations appear quite often. In this section we focus on regular ones though, so let us give some regular examples:

c) A nowhere vanishing vector field X is an action of $\mathbb{R}$ on M; its $C_c^\infty(M)$-linear span is a regular foliation.

d) A special case of the previous example is the Kronecker flow on the torus. Just consider T^2 as a quotient of $\mathbb{R}^2$ and put $X = \dfrac{d}{dx} + \theta\dfrac{d}{dy}$; If $\theta \in \mathbb{Q}$ then leaves are circles, while if θ is irrational then $\mathbb{R}$ is injected as a dense leaf.

e) **Horocyclic foliation:** Let Γ be a cocompact subgroup of $SL(2,\mathbb{R})$ and put $M = SL(2,\mathbb{R})/\Gamma$. The real line is embedded in $SL(2,\mathbb{R})$ by $t \mapsto \begin{pmatrix} 1 & 0 \\ t & 1 \end{pmatrix}$, therefore $\mathbb{R}$ acts on M. It turns out that this action is ergodic and there exist dense leaves.

1.2   Holonomy groupoid

Given a regular foliation $(M, \mathcal{F})$, a very useful model for the (often quite singular) space of leaves $M/\mathcal{F}$ is the holonomy groupoid. We give here an overview of this object in the regular case, as it will be the main tool for various constructions throughout this sequel.

The idea is that we wish to put a smooth structure on the equivalence relation of "belonging in the same leaf"

$$\{(x, x') \in M \times M | L_x = L_{x'}\}$$

Thinking about the dimension of this manifold, and in view of 1.1, we see that there are $p+q$ degrees of freedom for x. Once x is chosen, there are q degrees of freedom for x'. Whence, a neighbourhood of (x, x'), where $x \in U \times T$ and $x' \in U' \times T'$ should be of the form $U \times U' \times T$. Therefore we need an identification of T with T':

Definition 1.4. A (small) holonomy of $(M, \mathcal{F})$ is a diffeomorphism $h : T \to T'$ such that $t, h(t)$ lie in the same leaf for all $t \in T$.

Examples are the identity and the map h defined by the transition functions in 1.1. It is easy to see that the collection of (small) holonomies is a pseudogroup. In general (small) holonomies identify points in nearby (small) transversals which lie in the same leaf.

We can also identify points which lie in transversals far away from each other using the connectivity of the leaves: Consider a smooth path $\gamma : [0, 1] \to M$ which lies entirely in a leaf L. Cover γ by foliation charts $W_i = U_i \times T_i$, $1 \leq i \leq n$.

Definition 1.5. The holonomy of the path γ is the germ of $h(\gamma) = h_{W_n, W_{n-1}} \circ \ldots \circ h_{W_2, W_1}$.

The holonomy groupoid of $(M, \mathcal{F})$ is $H(\mathcal{F}) = \{(x, x', h(\gamma))\}$ where x, x' lie in the same leaf and γ is a path in the leaf with $\gamma(0) = x$ and $\gamma(1) = x'$.

A manifold[3] chart at $(x, x', h(\gamma)) \in H(\mathcal{F})$ is constructed like this: Let $U \times T$ a chart of M at x and $U' \times T'$ a chart at x'. Identifying T and T' with $h(\gamma)$, a chart of $(x, x', h(\gamma))$ is $U \times U' \times T$, whence $H(\mathcal{F})$ has dimension $2p + q$.

Moreover, $H(\mathcal{F})$ carries a natural groupoid structure with $t(x, x', h(\gamma))) = x$, $s(x, x', h(\gamma)) = x'$ and $(x, x', h(\gamma)) \circ (x', y, h(\delta)) = (x, y, h(\gamma) \circ h(\delta))$.

In fact $H(\mathcal{F})$ is the smallest[4] possible Lie groupoid whose Lie algeboid is F. Two natural bigger groupoids with algebroid F are:

- $G_1 = \{(x, y, h) : x, y \in L\}$ where h is any holonomy such that $h(x) = y$.

- $G_2 = \{(x, y, [\gamma]) : x, y \in L\}$ where $[\gamma]$ is the homotopy class of the path γ connecting x and y.

The following properties may be found in various places in the literature, *e.g.* in [13].

Proposition 1.6.   a) *The holonomy groupoid is a subgroupoid of G_1 and a quotient of G_2.*

  b) *The holonomy groupoid is the s-connected component of G_1. It is therefore an open subgroupoid of G_1. Moreover, it is dense in G_1.*

  c) *The set of leaves L such that the holonomy group $H(\mathcal{F})_x^x = s^{-1}(x) \cap t^{-1}(x)$ vanishes at every $x \in L$ (namely the set of leaves with no holonomy) is a dense G_d subset of M.*

1.3   The Laplacian(s) of the Kronecker foliation

Let us start with the example of the Kronecker flow in 1.3 d), assuming that the slope is irrational. Namely, the module $\mathcal{F}$ is the one spanned by the vector field $X = \dfrac{d}{dx} + \theta \dfrac{d}{dy}$ where $\theta \in \mathbb{R} \setminus \mathbb{Q}$. There are two Laplacians involved here:

$$\Delta_L = -\frac{d^2}{dx^2}, \text{ acting on } L^2(\mathbb{R}); \quad \Delta_{T^2} = -X^2, \text{ acting on } L^2(T^2)[5].$$

By Fourier transform, Δ_L becomes mupltiplication by ξ^2 of functions in $L^2(\mathbb{R})$. Its spectrum is an interval $[0, +\infty)$. Applying the Fourier transform to Δ_{T^2} we get a multiplication operator by $(n + \theta k)^2$, acting on $L^2(\mathbb{Z}^2)$. It turns out that the spectrum of this operator is *dense* in $[0, +\infty)$. Whence, we obtain two Laplacians with the same spectrum, albeit their spectrum is not of the same nature: The spectrum of Δ_L is Lebesgue, while the spectrum of Δ_{T^2} is discrete. Nevertheless, the following questions arise naturally:

- To what extend do the spectra of Δ_L and Δ_M coincide for any regular foliation $(M, \mathcal{F})$?

- If so, how can we calculate this spectrum?

---

[3]The holonomy groupoid of a regular foliation always carries a manifold structure (finite-dimensional) albeit its topology is not always Hausdorff. One example of this is the Reeb foliation on the solid torus.

[4]This means that if G is a Lie groupoid with connected s-fibers which integrates F, there exists an onto morphism of Lie groupoids $G \to H(\mathcal{F})$.

[5]In order to define the space of square-integrable functions on a manifold M, one usually needs to specify a measure on M. The way around such a choice is to work with sections of the bundle of 1-densities associated with M instead of general functions. This way we obtain an intrinsic definition of $L^2(M)$, free from any choices.

- What about singular foliations?

The regular case was dealt with by A. Connes and the material in [1], [2] and [3] generalizes Connes' method to the singular case. In the next section we make a start explaining the work of Connes in the regular case.

2 Spectral theory of Laplacians in the Regular case

The answer to the first question uses Connes' construction of the C^*-algebra $C^*(\mathcal{F})$ associated naturally with the foliation $(M, \mathcal{F})$. It turns out that $L^2(M)$ and $L^2(L)$ are natural faithful representations of $C^*(\mathcal{F})$ and the Laplacians Δ_M and Δ_L arise from a certain unbounded multiplier Δ of $C^*(\mathcal{F})$. Let us give an overview of these constructions.

2.1 The foliation C^*-algebra

Using the holonomy groupoid $H(\mathcal{F})$ one may associate a C^*-algebra to a regular foliation $(M, \mathcal{F})$. It is an instance of the more general construction of a C^*-algebra for any[6] Lie groupoid G over M, given by Connes in [7], which we describe here briefly:

The space $C_c^\infty(G)$ becomes a $*$-algebra with

- involution $f^*(g) = \overline{f(g^{-1})}$;

- convolution defined by a formula $f*h(g) = \displaystyle\int_{g_1 g_2 = g} f(g_1)h(g_2)$, where integration is understood along the fibers of the composition map $G \times_{s,t} G \to G$; such an integration can be defined using Haar systems[7]

The completion of $C_c^\infty(G)$ is possible due to the following results proven in [14]:

Proposition 2.1. a) *The (continuous) $*$-representations of the $*$-algebra $C_c^\infty(G)$ are in one to one correspondence with unitary representations of the groupoid[8] G.*

b) *An L^1-estimate shows that, for $f \in C_c^\infty(G)$ the map $f \mapsto \sup_\pi \|\pi(f)\|$ (where the supremum is over all such representations) is bounded. This estimate is*

$$\|f\|_1 = \sup_{u \in M} max\{\int_{G^u} |f(x)|d\lambda^u(x), \int_{G_u} |f(x)|d\lambda_u(x)\}$$

For every $u \in M$ we have the left-regular representation ρ of $C_c^\infty(G)$ on $L^2(G_u)$, given by the same formula as the convolution. This gives rise to two completions:

---

[6]The groupoid G in our account is supposed to be Hausdorff. However, Connes gave a construction of groupoid C^*-algebra(s) for non-Hausdorff Lie groupoids as well.

[7]A Haar system is (roughly) a smooth family $\lambda = (\lambda_u)_{u \in M}$ of invariant measures on the s-fibers $G_u = s^{-1}(u)$ of the groupoid. For a precise definition see *e.g.* [14] or half-densities. The inversion map $\iota : G \to G$ allows us to define a corresponding family $(\lambda^u)_{u \in M}$ of invariand measures on the t-fibers $G^u = t^{-1}(u)$.

[8]We elaborate on these representations in the Appendix.

Definition 2.2.   a) The full C^*-algebra $C^*(G)$ is the completion of $C_c^\infty(G)$ with respect to the norm $f \mapsto \sup_\pi \|\pi(f)\|$. (Here the supremum is taken over all possible representations π.)

  b) The reduced C^*-algebra $C_r^*(G)$ is the completion of $C_c^\infty(G)$ with respect to the norm $f \mapsto \sup_{u \in M} \|\rho_u(f)\|$.

2.2   Longitudinal Pseudodifferential calculus

We will also need to realize the foliation C^*-algebra by specific operators. Notice that the Lie algebra $\mathcal{F}$ of vector fields tangent to the foliation acts on $C_c^\infty(G)$ by unbounded multipliers. The algebra generated is the algebra of differential operators. Using the Fourier transform we can write such an operator acting by left multiplication on $f \in C_c^\infty(G)$ as:

$$P(f)(x,y) = \int exp(i < \phi(x,z), \xi >) \alpha(x,\xi) \chi(x,z) f(z,y) d\xi dz \tag{2.1}$$

where

- ϕ is the *phase* function: through a local diffeomorphism (tubular neighbourhood map) defined on an open subset $\widetilde{\Omega} \sim U \times U \times T \subset H(\mathcal{F})$ (where $\Omega = U \times T$ is a foliation chart), ϕ is given by $\phi(x,z) = x - z \in F_x$

- χ is a cut-off function; namely a smooth function on G such that $\chi(x,x) = 1$ on a compact subset of Ω and $\chi(x,z) = 0$ for $(x,z) \neq \widetilde{\Omega}$

- $\alpha \in C^\infty(F^*)$ is a polynomial on ξ. It is called the *symbol* of P.

We can make sense of (2.1) for more general symbols, in particular *poly-homogeneous* ones: those whose asymptotic expansion is $\alpha(u, \xi) \sim \sum_{k \in \mathbb{N}} \alpha_{m-k}(u, \xi)$ where α_j is homogeneous of degree j outside a neighbourhood of $M \subset F^*$. Note that α_m is called the *principal symbol* These are the *pseudodifferential operators* of order $m \in \mathbb{Z}$. The next result [17] shows the role of the foliation C^*-algebra:

Proposition 2.3.   a) *Negative order pseudodifferential operators are elements of* $C^*(\mathcal{F})$.

  b) *zero order pseudodifferential operators are multipliers of* $C^*(\mathcal{F})$.

  c) *Let* $\Psi^*(\mathcal{F})$ *be the closure of the zero order pseudodifferential operators in the multipliers algebra. The multiplicativity of the principal symbol map* σ *induces a short exact sequence of* C^*-*algebras:*

$$0 \to C^*(\mathcal{F}) \to \Psi^*(\mathcal{F}) \xrightarrow{\sigma} C(SF^*) \to 0 \tag{2.2}$$

(where SF^* *is the sphere bundle of* F^*. *)*

Example 2.4. Let M be a compact manifold and $\mathcal{F} = C^\infty(M, F)$ a foliation on M. Choose a basis of sections $X_1, \ldots, X_n \in \mathcal{F}$ and consider the Laplacian

$$\Delta = \sum_{i=1}^n X_i^* X_i$$

This is an unbounded multiplier of $C^*(\mathcal{F})$. In fact, viewing $C^*(\mathcal{F})$ as a Hilbert module, Baaj [6] and Woronowicz [18] showed that Δ is a *regular multiplier*, which means:

a) Δ is densely defined and closed; that is to say, the graph of Δ is a closed (right) submodule of $C^*(\mathcal{F}) \times C^*(\mathcal{F})$;

b) Δ has a densely defined (closed) adjoint Δ^*;

c) $graph\Delta \oplus (graph\Delta)^\perp = C^*(\mathcal{F}) \times C^*(\mathcal{F})$

More generally, in [17] it was shown that:

Proposition 2.5. *Elliptic pseudodifferential operators of positive order are regular unbounded multipliers of $C^*(\mathcal{F})$ (in the previous sense).*

2.3 Laplacians revisited

Now let us focus on the case of a regular foliation $\mathcal{F}$ on a compact manifold M. Consider the Laplacian $\Delta = \sum_{i=1}^{p} X_i^* X_i$ defined in 2.4. One sees easily that $L^2(M)$ and $L^2(L)$ are representations of $C^*(\mathcal{F})$ and $C_r^*(\mathcal{F})$. Baaj and Woronowicz proved the following:

Proposition 2.6. *Every representation of a Hilbert module extends to regular multipliers.*

This way the Laplacian Δ gives rise to Laplacians Δ_L acting on $L^2(L)$ and Δ_M acting on $L^2(M)$.

Theorem 2.7. *The Laplacians Δ_L and Δ_M are essentially self-adjoint.*

Proof. It follows directly from 2.6 and the simple fact the extension of a representation of Hilbert modules to regular multipliers preserves the adjoints. □

Theorem 2.8. *If all the leaves are dense and the holonomy groupoid is amenable[9] then Δ_L and Δ_M have the same spectrum.*

Proof. T. Fack and G. Skandalis in [12] showed that if all leaves are dense then the foliation C^*-algebra is simple. Whence all its representations are faithful. We conclude because every injective morphism of C^*-algebras is isometric and isospectral. □

2.4 Calculating the spectrum

Connes showed that in many cases it is possible to predict the possible gaps in the spectrum of the longitudinal Laplacian (Δ_M or Δ_L). His method uses the fact that gaps in the spectrum correspond to projections of $C^*(\mathcal{F})$, a fact that naturally leads to the use of K-theory. Let us sketch out an illustrating example:

Let Γ be a discrete co-compact subgroup of $SL(2, \mathbb{R})$. Much like example 1.3 e), the "$ax + b$" group acts on the compact manifold $M = SL(2/R)/\Gamma$. Leaves are the orbits of the "$x + b$" group. The associated holonomy groupoid turns out to be amenable.

Now the Haar measure on $SL(2, \mathbb{R})$ induces a measure on M which is invariant by the action of the "$ax + b$" group. Invariance by the "$x + b$"-subgroup gives rise to a trace on $C^*(\mathcal{F})$ which is faithful since $C^*(\mathcal{F})$ is simple.

---

[9]Amenability here means that the full and the reduced foliation C^*-algebras coincide.

On the other hand, the "ax"-subgroup induces an action of $\mathbb{R}_+^*$ on $C^*(\mathcal{F})$ which scales the trace. Whence, the image of $K_0(C^*(\mathcal{F}))$ is a countable subgroup of $\mathbb{R}$, invariant under the $\mathbb{R}_+^*$ action. Since this is satisfied only by $\{0\}$, it follows that $K_0(C^*(\mathcal{F}))$ vanishes, whence $C^*(\mathcal{F})$ has no projections. Therefore:

Proposition 2.9. *Let Γ be a discrete co-compact subgroup of $SL(2, \mathbb{R})$. Consider the regular foliation by the orbits of the "$ax+b$"-group action on $M = SL(2/R)/\Gamma$. The spectrum of the longitudinal Laplacian is an interval.*

Remark 2.10. Note that Connes' method to calculate the spectrum relies on the undersntanding of the K-theory $K_0(C^*(\mathcal{F}))$. Recall that the Baum-Connes assembly map predicts this K-theory group.

3 Singular foliations

In this section we are discussing singular foliations. We start with an overview of the work of C. Debord [10], [9], and then pass to the general case of a singular foliation.

3.1 Almost regular foliations

As we discussed in 1.2, a regular foliation $\mathcal{F}$ on a compact manifold M may be thought of as a projective $C^\infty(M)$-submodule of the vector fields of M, closed by the Lie bracket. However, such modules also appear for foliations with singularities; For instance, put $\mathcal{F}$ the span of a single vector field $X \in \mathcal{X}(M)$ such that the interior of $\{x \in M : X(x) = 0\}$ is empty.

In fact, the previous example is a special case of the following case studied by Debord:

Definition 3.1. An almost regular foliation $\mathcal{F}$ is a finitely generated and projective $C^\infty(M)$-submodule of $\mathcal{X}(M)$, closed by the Lie bracket.

Equivalently, $\mathcal{F} = \rho(\Gamma A)$ iwhere A is a Lie algebroid A over M, such that for every $x \in M$, its anchor map $\rho_x : A_x \to T_xM$ is injective in a dense subset of M. Debord and independently Crainic and Fernandes [8][10] showed:

Theorem 3.2. *Every almost injective Lie algebroid A is integrable. Namely, there is a Lie groupoid G with connected s-fibers whose Lie algebroid is A.*

Since there is always a Lie groupoid around in the almost regular case, we automatically obtain all the apparatus we discussed in the regular case. In particular, theorems 2.7 and 2.8 hold verbatim in this case, and their proof is exactly the same.

As for the calculation of the Laplacian's spectrum, we obtain that there are no gaps for a foliation defined by the "$ax + b$"-group action on manifold with conic singularities using a finite covolume subgroup of $SL(2, \mathbb{R})$.

---

[10]Actually Debord constructed the smallest groupoid that integrates an almost injective Lie algebroid A. For this reason it is legitimate to call this the *holonomy* groupoid of $\mathcal{F}$. On the other hand, Crainic and Fernandes proved that any almost injective Lie algebroid is integrable. The Lie groupoid they constructed is the biggest one that integrates A. It is connected and s-simply connected.

3.2 Stefan-Sussmann foliations

Dropping the projectivity of the module we obtain the following definition:

Definition 3.3. Let M be a compact manifold. A singular foliation is a finitely generated $C^\infty(M)$-submodule of vector fields, stable under the Lie bracket.

It was shown by Stefan [15] and Sussmann [16] that such sub-modules integrate to a partition of M to immersed submanifolds with non-constant dimension.

Remark 3.4. Let L be a leaf at $x \in M$. The tangent space $T_x L$ is the analogue of the fiber F_x in the regular case. However, in the singular case there is a second fiber $\mathcal{F}_x$ defined as the quotient of $\mathcal{F}/I_x\mathcal{F}$, where $I_x\mathcal{F}$ is the (maximal) ideal of $\mathcal{F}$ generated by $f \cdot X$, where $X \in \mathcal{F}$ and $f \in C^\infty(M)$ is a function that vanishes at x. This is a finite-dimensional vector space since $\mathcal{F}$ is finitely generated. The evaluation map gives rise to a short exact sequence

$$0 \to \mathfrak{g}_x \to \mathcal{F}_x \xrightarrow{ev_x} T_x L \to 0$$

Whence, the dimension of $\mathcal{F}_x$ as x runs through M is upper semi-continuous. The kernel $\mathfrak{g}_x$ is a Lie algebra and it vanishes on leaves of maximal dimension. (A detailed account of all this is in [4].)

A fundamental difference of singular foliations (even almost regular ones) from regular ones is that the partition to leaves no longer determines the dynamics. In other words, there can be more than one modules $\mathcal{F}$ which integrate to the same partition to leaves. This is better illustrated in the next examples:

Examples 3.5.   a) Consider $\mathbb{R}$ foliated by $(-\infty, 0), \{0\}, (0, +\infty)$. The module $\mathcal{F}$ may be the one generated by $x^n \dfrac{\partial}{\partial x}$, for any $n \in \mathbb{N}$. For every n this is an almost regular foliation. Moreover, there is a preferred choice of module, say the one corresponding to $n = 1$.

b) Consider the partition of $\mathbb{R}^2$ to $\{(0,0)\}$ and $\mathbb{R}^2 \setminus \{(0,0)\}$. These are the orbits of the action of $GL(2, \mathbb{R})$ on $\mathbb{R}^2$, or the action of $SL(2, \mathbb{R})$, or $\mathbb{C}^*$ (and $\mathcal{F}$ the image of the corresponding infinitesimal action). If we consider the action of $\mathbb{C}^*$ we get a projective module, whence an almost regular foliation. However, if we consider the action of $GL(2, \mathbb{R})$ or $SL(2, \mathbb{R})$ the associated module is no longer projective and we get a Stefan-Sussmann foliation.

In [1], [2], [3] the following were given in this more general setting:

- A holonomy groupoid.

- The foliation C^*-algebra(s) and its representation theory.

- The cotangent "bundle": It is not really a bundle, since the dimension of the fibers is not constant. However, it is a nice locally compact space.

- The longitudinal pseudodifferential calculus, together with the following results:

  a) The extension of C^*-algebras

  $$0 \to C^*(\mathcal{F}) \to \Psi^*(\mathcal{F}) \to C_0(\mathcal{F}^*) \to 0$$

  b) Elliptic operators of positive order are regular unbounded multipliers.

All of these constructions are highly technical. That is because of the difficulty with the topology of the holonomy groupoid when $\mathcal{F}$ is a singular foliation. We elaborate on this in the next section:

3.3 The holonomy groupoid in the singular case

Recall that for a singular foliation the partition to leaves may arise from different modules $\mathcal{F}$ of vector fields. This forces us to choose the dynamics of the foliation *a priori*, so the holonomy groupoid is expected to keep track of this choice rather than the partition to leaves. This suggests that the construction of $H(\mathcal{F})$ using path holonomies (in §1.1) is no longer sufficient in the singular case: We need a way to keep track of the path-holonomies determined by the original choice of vector fields that define the foliation. This is achieved with the notion of a *bi-submersion*. Let us give an overview of this:

Let $x \in M$ and consider the fiber $\mathcal{F}_x$. A basis of this linear space determines a finite number of vector fields $X_1, \ldots, X_n \in \mathcal{F}$ which generate $\mathcal{F}$ in a neighborhood of the point x. Let $U \subseteq \mathbb{R}^n \times M$ a neighborhood of $(0, x)$ where the exponential map

$$t(\lambda_1, \ldots, \lambda_n, y) = exp(\sum_{i=1}^{n} \lambda_i X_i)$$

is defined. The map $t : U \to M$ is a submersion, as well as the projection $s : U \to M$. By definition, for every $y \in s(U)$, the map t sends the fiber $s^{-1}(y)$ to the leaf L_y. In other words, the space U carries the (singular) foliation

$$s^{-1}(\mathcal{F}) = t^{-1}(\mathcal{F}) = C_c^{\infty}(U; \ker ds) + C_c^{\infty}(U; \ker dt) \tag{3.1}$$

Definition 3.6. Let $(M, \mathcal{F})$ be a (singular) foliation. A bi-submersion is a triple (U, t, s), where U is a smooth manifold, $s, t : U \to M$ are submersions and (3.1) is satisfied.

The role of bi-submersions is that they provide a very stable way to record those leaf-preserving local diffeomorphisms near the identity which arise from the original choice of module $\mathcal{F}$. More precisely, let $\mathbf{b}$ be a *bisection* of (U, t, s), namely a closed submanifold of U where the restriction of s and t are local diffeomorphisms. Then the map $\phi_{\mathbf{b}} : t \circ (s \mid_{\mathbf{b}})^{-1}$ is a leaf-preserving local diffeomorphism. It is "near the identity" in the following sense: Given a neighbrourhood V of x in M, the bisection $\mathbf{b}_0 = \{(0, \ldots, 0, y) : y \in V\}$ carries the identity, namely $\phi_{\mathbf{b}_0} = id_V$.

The holonomy groupoid $H(\mathcal{F})$ is a quotient of the collection of all bi-submersions we may construct over M. It has a quotient topology, which is very pathological. However, its open sets have smooth covers. That is to say that for every bi-submersion (U, t, s) the quotient map $\sharp : U \to H(\mathcal{F})$ is a smooth cover an open subset in $H(\mathcal{F})$. Let us give some examples to illustrate the holonomy groupoid in the singular case.

Examples 3.7.   a) In the (almost) regular case $H(\mathcal{F})$ coincides with the groupoid constructed by Debord (see [1]).

b) Let G be a Lie groupoid and $\mathcal{F}$ the image by the anchor map of the sections of its Lie algebroid AG. Then G is by definition a bi-submersion. The holonomy groupoid is a quotient of G.

c) Let X be a vector field on M which has non-periodic integral curves in a neighborhood of the boundary of its vanishing set $\{X = 0\}$. Put $\mathcal{F}$ the module spanned by X. The holonomy groupoid (see [4]) has three components:

$$H(\mathcal{F}) = H(X) \mid_{\{X \neq 0\}} \cup Int(X = 0) \cup \mathbb{R} \times \partial\{X = 0\}$$

d) Consider the module $\mathcal{F}$ defined by the (infinitesimal) action of $SL(2,\mathbb{R})$ on $\mathbb{R}^2$. In [1] it was shown that

$$H(\mathcal{F}) = (\mathbb{R}^2 \setminus \{0\})^2 \cup SL(2,\mathbb{R}) \times \{0\}$$

The topology of this groupoid is so bad that for every $x \in \mathbb{R}^2 \setminus \{0\}$ the sequence $(\frac{x}{n}, \frac{x}{n}) \in H(\mathcal{F})$ converges to every element of the stabilizer group (i.e. to every element of the real line)!

The holonomy groupoid of a singular foliation is rarely smooth, however Debord in [11] proved the following:

Proposition 3.8. *Given any singular foliation $\mathcal{F}$ on M, the s-fibers of the holonomy groupoid $H(\mathcal{F})$ are smooth manifolds.*

Note that in [4] the smoothness of the s-fibers is linked with the Crainic-Fernandes integrability obstructions [8].

4   Longitudinal Laplacians for Singular Foliations

Bi-submersions make possible the generalization of Connes' construction[11] of the C^*-algebra of a non-Hausdorff groupoid to the holonomy groupoid of a singular foliation $(M, \mathcal{F})$. Very roughly, if $\{(U_i, t_i, s_i)\}_{i \in I}$ is a collection of bi-submersions such that $\cup_{i \in I} s_i(U_i) = M$ then it turns out that a quotient $\mathcal{A}$ of $\oplus_{i \in I} C_c^\infty(U_i)$ admits an involution and a convolution. Similarly with the regular case, $\mathcal{A}$ can be completed to the foliation C^*-algebra(s). The following was proven in [2]

Proposition 4.1. *Let M be a smooth compact manifold and $X_1, \ldots, X_n$ smooth vector fields such that $[X_i, X_j] = \sum_{k=1}^{n} f_{ij}^k X_k$ for some $f_{ij}^k \in C^\infty(M)$. Then $\Delta = \sum_{i=1}^{n} X_i^* X_i$ is a regular unbounded mupltiplier of $C_r^*(\mathcal{F})$.*

It follows that Δ induces essentially self-adjoint operators in every representation, particularly in $L^2(M)$ and $L^2(L)$.

Theorem 4.2. *Assume that the dense open set $\Omega \subset M$ where leaves have maximal dimension has Lebesgue measure 1. Assume the restriction of all leaves to Ω are dense in Ω. Assume that the holonomy groupoid of the restriction of $\mathcal{F}$ to Ω is Hausdorff and amenable. Then Δ_M and Δ_L have the same spectrum.*

Proof. By [12] it follows that $C^*(\Omega, \mathcal{F})$ is simple and sits as a two-sided ideal in $C^*(\mathcal{F})$. So $L^2(L)$ and $L^2(M)$ are faithful representations of $C^*(\Omega, \mathcal{F})$, whence they are weakly equivalent. We conclude because the extensions of these representations to multipliers are also weakly equivalent. □

---

[11]Actually the construction of the foliation C^*-algebra(s) given in [1] is inspired by Connes' construction of a groupoid C^*-algebra for a non-Hausdorff Lie groupoid.

5 The Baum-Connes conjecture and the calculation of the spectrum

Recall that Connes' calculation of the spectrum in the regular case requires to know the "shape" of $K_0(C^*(\mathcal{F}))$, and this leads to the Baum-Connes assembly map.

In the singular case, it turns out that the union of leaves of a given dimension are locally closed subsets and they give rise to a decomposition series for the foliation C^*-algebra. This leads to the following conjecture:

Conjecture 5.1. *Given a singular foliation $(M, \mathcal{F})$, the Baum-Connes assembly map is an isomorphism if and only if it is an isomorphism for every stratum of leaves of equal dimension.*

Of course, the difficulty here is to actually write down the Baum-Connes assembly map; to this end one needs to understand the classifying space of proper actions of the (very singular) holonomy groupoid. Also, in general the Baum-Connes assembly map is a kind of an analytic index map; note that in [3] the analytic index map was constructed for any singular foliation.

A Groupoid representations on Hilbert bundles

Let G be a Lie groupoid over a manifold M with a Haar system $\lambda = (\lambda^x)_{x \in M}$. Let μ be a measure on M. Define a measure $\mu \circ \lambda$ on G by $\mu \circ \lambda(f) = \int_M (\lambda^x(f)) d\mu(x)$. Using the inversion map $g \mapsto g^{-1}$ we define similarly another measure $\widetilde{\mu \circ \lambda}$.

Definition A.1. The measure μ is quasi-invariant if $\mu \circ \lambda$ is equivalent to $\widetilde{\mu \circ \lambda}$.

We denote δ the associated Radon-Nikodym derivative.

Definition A.2. A representation of G is a triple $(\mu, \mathcal{H}, \theta)$ where

- μ is a quasi-invariant measure on M;

- $\mathcal{H} = (H_x)_{x \in M}$ is a measurable field of Hilbert spaces;

- θ is a unitary representation of G on $\mathcal{H}$, namely:

  a) For $g \in G$, $\theta(x) : H_{s(g)} \to H_{t(g)}$ is a unitary operator.
  b) For (almost) all composable (g_1, g_2) we have $\theta(g_1 g_2) = \theta(g_1)\theta(g_2)$.
  c) The field $(\theta(g))_{g \in G}$ is measurable.

Every such representation θ induces a representation of $C_c^\infty(G)$ on $\mathbb{H} = \int_M^\oplus H_x d\mu(x)$ by

$$(\pi_\theta(f)(\xi))(g) = \int f(g)\delta(g)^{\frac{1}{2}}\theta(g)\xi(s(g))d\lambda^x(g)$$

for every $f \in C_c^\infty(G)$, $\xi \in \mathbb{H}$ and $g \in G$. The following was proven in [14] and [12].

Theorem A.3. *Every representation of $C_c^\infty(G)$ is the integrated form of a representation of the groupoid G.*

References

[1] I. Androulidakis and G. Skandalis. The holonomy groupoid of a singular foliation. *J. Reine Angew. Math.*, 626:1–37, 2009.

[2] I. Androulidakis and G. Skandalis. Pseudodifferential calculus on a singular foliation. *J. Noncommut. Geom.*, 5(1):125–152, 2011.

[3] I. Androulidakis and G. Skandalis. The analytic index of elliptic pseudodifferential operators on a singular foliation. *J. K-Theory*, 8(3):363–385, 2011.

[4] I. Androulidakis and M. Zambon. Smoothness of holonomy covers for singular foliations and essential isotropy. *Math. Z.*, 275(3-4):921–951, 2013.

[5] I. Androulidakis and M. Zambon. Holonomy transformations for singular foliations. *Adv. Math.*, 256:348–397, 2014.

[6] S. Baaj and P. Julg. Théorie bivariante de Kasparov et opérateurs non bornés dans les C^*-modules hilbertiens. *C. R. Acad. Sci., Paris, Sér. I*, 296:875–878, 1983.

[7] A. Connes. Sur la théorie non commutative de l'intégration. Algebres d'operateurs, Semin. Les Plans-sur-Bex 1978, Lect. Notes Math. 725, 19-143 (1979)., 1979.

[8] M. Crainic and R. L. Fernandes. Integrability of Lie brackets. *Ann. Math. (2)*, 157(2):575–620, 2003.

[9] C. Debord. Holonomy groupoids of singular foliations. *J. Differ. Geom.*, 58(3):467–500, 2001.

[10] C. Debord. Local integration of Lie algebroids. In *Lie algebroids and related topics in differential geometry. Proceedings of the conference, Warsaw, Poland, June 12–18, 2000*, pages 21–33. Warsaw: Polish Academy of Sciences, Institute of Mathematics, 2001.

[11] C. Debord. Longitudinal smoothness of the holonomy groupoid. *C. R., Math., Acad. Sci. Paris*, 351(15-16):613–616, 2013.

[12] T. Fack and G. Skandalis. Sur les représentations et ideaux de la C*-algèbre d'un feuilletage. *J. Oper. Theory*, 8:95–129, 1982.

[13] J. Pradines and B. Bigonnet. Graphe d'un feuilletage singulier. (Graph of a singular foliation). *C. R. Acad. Sci., Paris, Sér. I*, 300:439–442, 1985.

[14] J. Renault. A groupoid approach to C*-algebras. Lecture Notes in Mathematics. 793. Berlin-Heidelberg-New York: Springer-Verlag. III, 160 p. DM 21.50; $ 12.70 (1980)., 1980.

[15] P. Stefan. Accessible sets, orbits, and foliations with singularities. *Proc. Lond. Math. Soc. (3)*, 29:699–713, 1974.

[16] H. J. Sussmann. Orbits of families of vector fields and integrability of distributions. *Trans. Am. Math. Soc.*, 180:171–188, 1973.

[17] S. Vassout. Unbounded pseudodifferential calculus on Lie groupoids. *J. Funct. Anal.*, 236(1):161–200, 2006.

[18] S. Woronowicz. Unbounded elements affiliated with C^*-algebras and non-compact quantum groups. *Commun. Math. Phys.*, 136(2):399–432, 1991.