Scuola Internazionale Superiore di Studi Avanzati - Trieste

PhD Thesis

A.A. 2013-2014

Geodesics and horizontal-path spaces in sub-Riemannian geometry

Thesis submitted for the fulfillment of the PhD in Geometry

Candidate:
Alessandro Gentile

Supervisors:
Andrei Agrachev
Antonio Lerario

SISSA - Via Bonomea 265 - 34136 TRIESTE - ITALY

Contents

Abstract

We study the topology of the space Ω_p of *horizontal* paths between two points e (the origin) and p on a step-two Carnot group G:

$$\Omega_p = \{\gamma : I \to G \,|\, \gamma \text{ horizontal}, \, \gamma(0) = e, \, \gamma(1) = p\}.$$

As it turns out, Ω_p is homotopy equivalent to an infinite dimensional sphere for $p \neq e$ and in particular it is contractible. The *energy* function:

$$J : \Omega_p \to \mathbb{R}$$

is defined by $J(\gamma) = \frac{1}{2}\int_I \|\dot{\gamma}\|^2$; critical points of this function are sub-Riemannian geodesics between e and p. We study the asymptotic of the number of geodesics and the topology of the sublevel sets:

$$\Omega_p^s = \{\gamma \in \Omega_p \,|\, J(\gamma) \leq s\} \quad \text{as } s \to \infty.$$

Carnot groups are sub-Riemannian tangent spaces in the same way Euclidean spaces are Riemannian tangent spaces: in Euclidean spaces there is only one geodesic joining two points and the sublevels Ω_p^s of the energy in the space of paths are contractible for every value of s such that Ω_p^s is not empty. But in the sub-Riemannian case a completely new behavior is experienced for the generic *vertical* p (a point in the second layer of the group). In this case we show that J is a Morse-Bott function: geodesics appear in isolated *families* (critical manifolds), indexed by the critical values of the energy. Denoting by l the corank of the horizontal distribution on G, we prove that:

$$\text{Card}\{\text{Critical manifolds with energy less than } s\} = O(s)^l.$$

Despite this evidence, Morse-Bott inequalities $b(\Omega_p^s) = O(s)^l$ are far from being sharp and we show that the following stronger estimate holds:

$$b(\Omega_p^s) = O(s)^{l-1}.$$

Thus each single Betti number $b_i(\Omega_p^s)$ $(i > 0)$ becomes eventually zero as $s \to \infty$, but the sum of all of them can possibly increase as fast as $O(s)^{l-1}$. In the case $l = 2$ we show that indeed

$$b(\Omega_p^s) = \tau(p)s + o(s) \quad (l = 2).$$

The leading order coefficient $\tau(p)$ can be *analytically* computed using the structure constants of the Lie algebra of G and it is generically not 0.

Using a dilation procedure, reminiscent to the rescaling for Gromov-Hausdorff limits, we interpret these results as giving some local information on the geometry of G (e.g. we derive for $l = 2$ the rate of growth of the number of geodesics with bounded energy as p approaches e along a vertical direction). Moreover we are able to use the dilation procedure to obtain a homotopy equivalence between the horizontal path spaces with bounded energy of a general sub-Riemannian manifold and of the Carnot group providing its nilpotent approximation. From this result we can find a lower bound for the number of geodesics between two points depending on their distance:

$$\mathrm{Card}\left(\mathrm{Crit}_J \cap \Omega^c_{x,p_\epsilon}\right) \geq \mathrm{const}(x)\epsilon^{-1},$$

where the distance $d(x, p_\epsilon) = \epsilon$ and Ω^c_{x,p_ϵ} is the space of horizontal paths between x and p_ϵ with energy less than c.

Chapter 1

Introduction

1.1 The space of horizontal paths

The relation between geodesics (in the sense of critical points of the Energy, not just minimizers) and the space of paths joining two points have been widely investigated in Riemannian geometry: by means of Morse theory one can retrieve informations about the topology of the space of paths from the structure of the geodesics and viceversa (see for example the work of Serre [29]).

In sub-Riemannian geometry we have a distribution Δ on a smooth manifold M; such distribution is endowed with a smoothly varying scalar product, namely a *sub-Riemannian metric*. The same study can be carried out: we restrict to the space of *horizontal paths*, namely those paths whose velocity belongs to the distribution Δ almost everywhere. Since the sub-Riemannian metric is given, the Energy of a path can be defined and it is still possible to define what geodesics are, like in Riemannian geometry.

The constraints given on the velocities by the distribution Δ are supposed to be *completely non-holonomic*, meaning that iterated brackets of vector fields in Δ generate all the tangent bundle TM: their main property is that *every* two points $p, q \in M$ can be joined by horizontal curves (*Chow-Rashevskii Theorem*). This case is the exact opposite of holonomic constraints, where the closure with respect to the brackets forces the motion to be constrained on submanifolds, as stated by the Frobenius' Theorem.

In Riemannian geometry, the way that the topology of the space of paths is related to the structure of geodesics depends on the topology of the manifold M, but locally there is not so much to say: on a sufficient small neighborhood, every two points are connected by only one geodesic. On the contrary, in sub-Riemannian geometry there are non trivial properties of both geodesics and horizontal-path space, even if we take the end-points to be "infinitesimally" close. Thus we focus on properties of both the set of the geodesics and the space of horizontal paths, that depend on the sub-Riemannian local data only and does not depend on topological properties of the manifold M.

It is natural then to begin with local models of sub-Riemannian structures; in the same way Euclidean spaces are local models for Riemannian manifolds, Carnot groups are local models for sub-Riemannian manifolds, and they carry a very rich structure.

1.2 The Heisenberg group

In order to get introduced to the subject, we can study what happens on the Heisenberg group, which is the only 3-dimensional Carnot group and therefore the only local model for 3-dimensional sub-Riemannian manifolds.

The Heisenberg group can be realized as a copy of $\mathbb{R}^3$ with coordinates (x, y, z), and its Lie algebra is generated by the following left invariant vector fields:

$$X \doteq \frac{\partial}{\partial x} - \frac{1}{2} y \frac{\partial}{\partial z}, \quad Y \doteq \frac{\partial}{\partial y} + \frac{1}{2} x \frac{\partial}{\partial z}, \quad Z \doteq \frac{\partial}{\partial z}.$$

We define the rank 2 distribution $\Delta \doteq \mathrm{span}\{X, Y\}$ and we define the sub-Riemannian metric by letting X and Y be orthonormal. Since $[X, Y] = Z$ we easily see that Δ is non-holonomic; moreover, one bracket only is needed to generate all the tangent bundle and it follows that the given structure has *step 2*.

In order to study the geodesics and the topology of horizontal paths, we fix without loss of generality (since everything is left-invariant) the initial point to be the identity element $e = (0, 0, 0)$. The geodesics starting from the origin e have a nice geometrical description: their projection on the xy-plane are circles, while the z-coordinate equals the area of the circle spanned by the projection (see figure 1.1). As a limit case we have circles of infinite radius, namely straight lines on the xy-plane. Moreover the length of the geodesics is equal to the length of their projection on the xy-plane with Euclidean metric. Depending on the choice of the final point, we find three possible cases for the structure of the geodesics. If p belongs to the xy-plane there is only one geodesic, which is a straight line, like in the Euclidean setting. If p does not belong to the xy-plane but it is also out of the z axis, the properties of non-holonomic geometry come into light and we find more and more geodesics as the ratio $\frac{|z|}{|x|+|y|}$ increases (for a quantitative study see [21]); the different geodesics make more and more revolutions, but in order to remain bounded

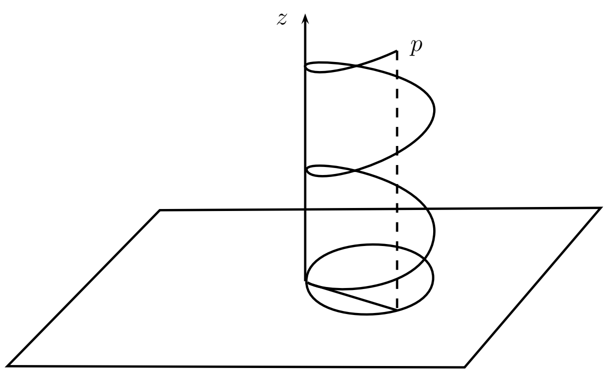

Figure 1.1: Geodesics in the Heisenberg group.

with the height, the radius of the circles has to get smaller and at some point we cannot reach the projection of p on the xy-plane anymore. Therefore, if the point p belongs to the z axis, we don't have any lower bound on the radius of the circles, so we have an infinite number of geodesics. More precisely: let $p = (0, 0, h)$ be the final point. To every geodesic we can associate the radius of the projection and the number of revolutions; let's say that the geodesic γ_n makes n revolutions and it has radius r_n. Since the height equals the area spanned by the projection, we must have

$$h = n \cdot \pi r_n^2;$$

we see that for every integer number n we can find a suitable radius in order to get the corresponding geodesics. Moreover, the Heisenberg structure is invariant under the group of the rotations of the xy-plane: applying a rotation of angle α to a geodoesic we find a new geodesic for every $\alpha \in S^1$.

The picture is the following: geodesics come in families parametrized by the integer numbers and each one of them is diffeomorphic to a circle S^1 in the space of horizontal curves. We will see that this picture is typical for step 2 Carnot groups with any possible corak l: there is a countable quantity of families of geodesics, and each family is diffeomorphic to a compact manifold in the space of horizontal paths; in the generic case the compact manifold is a product of circles.

1.3 The asymptotic behavior of the topology

Now we want to find the relation between the structure of the geodesics and the topology of the space of horizontal paths, denoted by Ω_p. To this aim the classical tool would be Morse theory, but in our case the critical points are not isolated. Though, they have a nice and quite regular structure: they are collected into compact manifolds. We will need then Morse-Bott theory, which generalizes Morse theory allowing the critical set to be the disjoint union of compact manifolds, called *non-degenerate critical manifolds*. We recall that Morse inequalities bound the Betti number of the manifold from above with the number of critical points. The most rough Morse inequality is obtained by summing over the integer that parametrizes the dimension of the homology cycles: we get that the total Betti number (the sum of all the Betti numbers) is bounded from above by the number of the critical points. In the Morse-Bott case, apparently we cannot do the same since the number of critical points is always infinite: but in the correct generalization of Morse inequalities the critical manifolds give their contribution with their total Betti number. We have then

$$b(M) \leq \sum_C b(C),$$

where C are the critical manifolds.

As usual in this kind of problems, instead of the length we will use another functional, namely the Energy. Since we consider horizontal curves and by definition of sub-Riemannian metric we can compute the norm of their velocities, we can define the length

$$\ell(\gamma) = \int_0^1 \|\dot{\gamma}\| dt,$$

and the Energy

$$J(\gamma) = \frac{1}{2} \int_0^1 \|\dot{\gamma}\|^2 dt,$$

which is easier to study but it produces the same critical points, the geodesics. This follows from the Schwarz inequality; moreover if γ is a geodesic we have $J(\gamma) = \frac{1}{2}\ell(\gamma)^2$. Notice that even though the horizontal curves come from non-autonomous ODEs, they are invariant under reparametrization: thus we have fixed the interval of definition to be $I = [0, 1]$ for the sake of exposition. Later we will choose the interval $[0, 2\pi]$ in order to have lighter computations.

Now we go back to the Heisenberg group and we study how the space of horizontal paths is built up as the union of the sublevels of the Energy. We expect a relation with the number of critical manifolds counted with their total Betti number, which is always 2 being the critical manifolds always homeomorphic to S^1. The geodesics γ_n have length $\ell(\gamma_n) = 2n\pi r_n$, thus their Energy is $J(\gamma_n) = \frac{1}{2} 4n^2\pi^2 r_n^2 = 2\pi n h$.

Let us call the sublevels of the Energy $\Omega_p^s \doteq \Omega_p \cap \{J \leq s\}$; since the Energy of the families of the geodesics grows linearly with the number of revolutions, we find a finite number of families with Energy less than s, where the number is (up to a fixed constant) the integer part of s. As we let the energy go to infinity, the number of families of geodesics explodes. Let's see what happens to the topology; it can be easily shown that Ω_p^s has the same homotopy type of a sphere, and that the dimension of the sphere is proportional to the level of the Energy s. The homotopy type changes with the Energy s, and converges to the homotopy type of the infinite dimensional sphere which is contractible.

We have seen two interesting facts: first, even though the total space of horizontal paths ending at p is contractible, the topology of the sublevels of the Energy changes with the value of the Energy but it is never trivial; only in the limit everything disappears. Second, the Morse-Bott inequalities are not sharp at all: while the total Betti number of the sublevels Ω_p^s is constantly equal to 2, the sum of the Betti numbers of the critical manifolds with Energy less than s grows proportional to s; it seems that the family of critical manifolds indexed by the integer numbers gives the same contribute to the topology of the manifold, independent on the number of critical manifolds we are considering.

1.4 On step 2 Carnot groups

Step 2 Carnot groups are the local models for the sub-Riemannian manifolds of step 2, namely those sub-Riemannian structures where the distribution Δ and the brackets of vector fields in the distribution span the whole tangent bundle. A step 2 Carnot group can be defined starting from its Lie algebra:

$$\mathfrak{g} = \Delta \oplus \Delta^2,$$

where Δ is the distribution and the brackets of left invariant vector fields in Δ give Δ^2, namely $[\Delta, \Delta] = \Delta^2$. Moreover $[\Delta^2, \mathfrak{g}] = 0$, thus the Lie algebra is nilpotent; the dimension of Δ^2 is the *corank* of the distribution Δ. We write explicitly the Lie bracket as

$$[X_i, X_j] = \sum_{k=1}^{l} a_{ij}^k Y_k,$$

and the skew-symmetric matrices $A_k \doteq (a_{ij}^k)$ will play an important role.

At first we identify the set of *vertical points*, analogous to the z axis in the Heisenberg group, where the most typical behavior of non-holonomic geometry is found. Vertical points are obtained by exponentiating Δ^2. Then we study how the topology of Ω_p^s changes when p is vertical, and at the same time we want to find the geodesics with Energy less than s. Since an exact computation is rather impossible, we focus on the order of growth of the above quantities with respect of the Energy by looking for upper bounds. We have the following result:

Theorem. *For the generic step 2 Carnot group G and the generic vertical point p, in the space of horizontal paths Ω_p we have*

$$\mathrm{Card}\{\textit{critical manifolds with energy less than } s\} = O(s)^l,$$

where l is the corank of the distribution.

This is still not enough in order to find what the right-hand side of the Morse inequalities would be; however we find what are the possible homeomorphism types of the critical manifolds:

Theorem. *For the generic step 2 Carnot group G and the generic vertical point p, the non-degenerate critical manifolds are homeomorphic to a product of spheres*

$$\underbrace{S^1 \times \cdots \times S^1}_{j \textit{ times}},$$

where $j \leq l$.

Since from the previous theorem we have that the total Betti number of the critical manifolds is bounded by a constant that depends only on the corank l, we can effectively bound the right-hand side of the Morse inequalities:

Theorem. *For the generic step 2 Carnot group G and the generic vertical point p, the total Betti number of the critical set of the energy J restricted to the sublevel Ω_p^s is bounded as follows:*

$$b\left(Crit_J \cap \Omega_p^s\right) \leq O(s)^l.$$

From the classical Morse inequalities we would find that $b(\Omega_p^s) \leq O(s)^l$; instead of doing this, we use Morse theory in order to find a homotopy equivalence between Ω_p^s and a finite dimensional submanifold of Ω_p^s that we can study as an intersection of l real quadrics. More precisely we prove that:

Theorem. *For the generic step 2 Carnot group G and the generic vertical point p, there is a constant c_p such that we have the following homotopy equivalence:*

$$\Omega_p^s \sim \partial\left(\Omega_p^s \cap L\right),$$

where $\dim L \leq c_p s$. *Moreover the manifold* $\partial\left(\Omega_p^s \cap L\right)$ *is an intersection of* l *quadrics on the unitary sphere in* L.

Given that, tools from semialgebraic geometry come into play: we use the results from [5] that allow to bound the total Betti number of intersections of quadrics on a sphere of given dimension to get an asymptotically finer bound for the total Betti number of Ω_p^s:

Theorem. *For the generic step 2 Carnot group* G *and the generic vertical point* p *we have*

$$b(\Omega_p^s) \leq O(s)^{l-1}. \tag{1.1}$$

In this case we have again only a bound from above for the total Betti number; but in the case of corank $l = 2$ we can do better. We find a formula that allows to compute analytically the leading coefficient of the asymptotic in terms of the structure constants of the Lie algebra, and such a leading coefficient is generically non-zero. We will show how the Carnot group data depend only on the linear space $W \subset \mathfrak{so}(d)$ generated by the matrices A_k that give the structure constants for the Lie algebra. This linear space can be identified with the space of covectors $(\Delta^2)^*$ via the map

$$\omega \mapsto \omega A \doteq \sum_{k=1}^{l} \omega_k A_k.$$

Every matrix ωA can be put in its canonical form

$$\omega A = \left(\alpha_1(\omega) J_2, \ldots, \alpha_m(\omega) J_2\right),$$

where J_2 is the standard symplectic matrix on $\mathbb{R}^2$; the basis is chosen in such a way that every $\alpha_j \geq 0$. Now we fix the ending point $p \in \Delta^2$, and we parametrize by $t \in R$ the line in $(\Delta^2)^*$ given by the equation $\langle \omega, p \rangle = 1$; the generalized eigenvalues α_j are seen as functions of t.

Theorem. *For the generic step 2 Carnot group* G *with corank* $l = 2$ *and the generic vertical point* p*, we have that*

$$b(\Omega_p^s) = \tau_p\, s + O(1),$$

where the leading coefficient τ_p *is greater than* 0 *and it is computed by the formula:*

$$\tau_p = \frac{1}{2} \int_{\mathbb{R}} \sum_{i=1}^{m} \left| \frac{d\alpha_i}{dt} \right| - \left| \sum_{i=1}^{m} \frac{d\alpha_i}{dt} \right|.$$

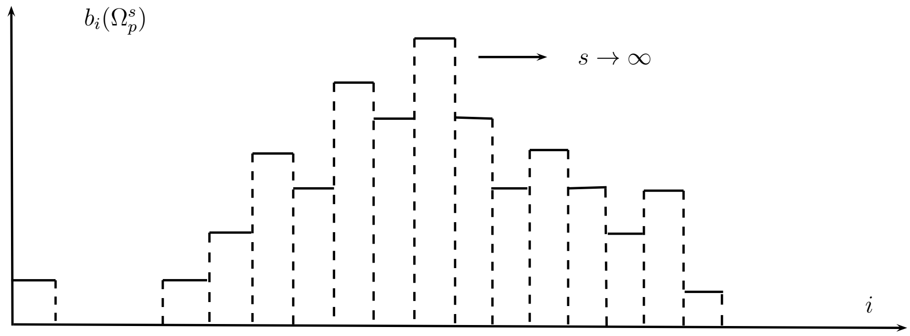

The previous theorem says that the total Betti number of the sublevel Ω_p^s tends to infinity as the energy increases: thus the (homological) complexity of the space of horizontal paths with bounded energy grows as the bound increases. Nevertheless every single Betti number (except the zero-th) becomes definitely zero and in the direct limit everything becomes trivial. If we look at the function $i \mapsto b_i(\Omega_p^s)$ and we let $s \to \infty$, what we see is something like a "wave" that moves towards infinity getting bigger and bigger.

1.5 General sub-Riemannian structures

Now we can go back to the problem on general sub-Riemannian manifolds. It has been proved in [20] that on a step 2 sub-Riemannian manifold the end-point map is a Serre fibration and we have

$$\pi_k(\Omega_{p,q}) \cong \pi_{k+1}(M),$$

where p, q are points in the sub-Riemannian manifold M and $\Omega_{p,q}$ is the space of horizontal paths between p and q. There is essentially no difference with the Riemannian case, since the topology of the horizontal-path spaces descends from the topology of the manifold M. Thus we focus on the study of the properties that depend on the sub-Riemannian structure, namely local properties of the distribution.

The key idea is that Carnot groups are first order approximations of general sub-Riemannian manifolds: it follows that we can pass locally the information retrieved in Carnot groups to the sub-Riemannian manifolds. Here locally means that we consider points "infinitesimally close" and the space of horizontal paths with an infinitesimal upper bound on the Energy.

Given a sub-Riemannian manifold M with distribution Δ, we fix a point $p \in M$ and we define F to be the End-point map, namely the map that associates to a horizontal path γ starting from p its final point $\gamma(1)$. We then fix a realization of the nilpotent approximation (the Carnot group that approximates the sub-Riemannian structure) in a neighborhood of p, namely a distribution that approximates the initial one and has good homogeneity properties. The Carnot group comes from the distribution since any choice of local vector fields that span the distribution generate the nilpotent Lie algebra associated to the Carnot group. This choice comes with a 1-parameter family of non isotropic dilations δ_ϵ that behaves well with respect to the sub-Riemannian distance.

We find the following weak homotopy equivalence for $\epsilon < \epsilon(s)$:

$$F^{-1}(\delta_\epsilon(p)) \cap \{J \leq \epsilon^2 s\} \sim \hat{F}^{-1}(p) \cap \{J \leq s\},$$

where $\hat{F}$ is the End-point map associated to the nilpotent approximation. The consequences of this homotopy equivalence are yet to be exploited; for instance from the study of the relative homology of different sublevels of the Energy we can find information about the number of "low Energy" geodesics between two close points, depending on their distance. More precisely, let us call $\nu(\Omega_q^s)$ the number of geodesics from p to q with Energy bounded by s: we have:

Theorem. *For the generic step 2 rank d distribution Δ on $\mathbb{R}^n$, and for almost every q in a small enough neighborhood of p (small enough in order to have privileged coordinates), for every $C > 0$ the quantity $\nu\left(\Omega^C_{\delta_\epsilon(q)}\right)$ goes to ∞ as $\epsilon \to 0$.*

Chapter 2

Background material

2.1 Sub-Riemannian geometry

2.1.1 Sub-Riemannian structures

For a more detailed exposition of sub-Riemannian geometry the reader is referred to [2]; for what concerns the present work we will not give the most general definition of a sub-Riemannian structure, since we are interested in the local study around regular points (regularity will be introduced later).

Definition 1 (sub-Riemannian structures)**.** *A* ***sub-Riemannian structure*** *on a smooth manifold M of dimension n is given by a distribution $\Delta \subset TM$, namely a collection of subspaces $\Delta_p \subset T_pM$ varying smoothly with respect to the point $p \in M$. We can assume that the* ***rank*** *d (dimension of Δ_p) is constant. Moreover, on the subspaces Δ_p we are given a scalar product, varying smoothly with respect to the point p as well; it is called* ***sub-Riemannian metric****.*

This definition generalizes Riemannian structures, where the distribution is just the whole tangent bundle. In the same way as the tangent bundle, a distribution can be given a local trivialization: given p in the manifold M with a rank d distribution Δ, we can find vector fields $\xi_1, \dots, \xi_d$ on a neighborhood U of p such that $\Delta_q = \text{span}(\xi_1(q), \dots, \xi_d(q))$ for every $q \in U$.

When such a structure is given, it is interesting to study the motion on the manifold M with the constraint on the velocity given by the distribution Δ.

Definition 2. *An absolutely continuous curve $\gamma : I \to M$, where I is a closed interval of $\mathbb{R}$, is said* ***horizontal*** *if its derivative (which exists almost everywhere) $\dot\gamma(t)$ belongs to the fiber $\Delta_{\gamma(t)}$.*

From now on we will consider horizontal curves only. Given the sub-Riemannian metric, we can compute the length of every horizontal curve since their velocity belongs to the distribution Δ by definition:

$$\ell(\gamma) = \int_I \|\dot\gamma(t)\| dt;$$

the length is independent of the parametrization.

The properties of the horizontal curves can be dramatically different depending on the local differential structure of the distribution Δ.

Given two vector fields X, Y in Δ on a neighborhood of a point $p \in M$, we consider their bracket $[X, Y]_p$ at p; now, for every couple of such vector fields we consider Δ_p together with the set of all possible brackets in p, which turns out to be a linear subspace of T_pM denoted by $[\Delta, \Delta]_p$ or Δ_p^2. By taking the union $\bigcup_{p\in M} \Delta_p^2$ we obtain the subset $\Delta^2 \subset TM$, which may not be a smooth sub-bundle of TM; the dimension of Δ_p^2 may also change with respect to p. The process of taking brackets can be repeated inductively: the brackets of all possible vector fields X in Δ and Y in Δ^m define the subset $\Delta^{m+1} \subset TM$. The subset Δ^m is called *m-th layer* of the distribution Δ. We can also take the union of the layers on $m \in \mathbb{N}$, namely $\text{Lie}(\Delta) = \bigcup_{m\in\mathbb{N}} \Delta^m$. The number σ of the layers is called *step* of the distribution Δ; some information about the layers can be summarized in a finite sequence of non-decreasing integer number, the *growth vector* defined as $(\text{dim}\Delta, \dots, \text{dim}\Delta^\sigma)$.

A distribution Δ is called *integrable* if $\text{Lie}\,\Delta = \Delta$: from the *Frobenius' theorem* it follows that for every point $p \in M$, there exists a smooth submanifold $N \subset M$ passing through p such that $\Delta = TN$. This means that in the case of integrable constraints on the velocities, the study of horizontal curves can be reduced to the Riemannian case on the manifold N.

However, the object of study in sub-Riemannian geometry is the very opposite case: the distribution Δ satisfies the *Hörmander condition* if $\text{Lie}\,\Delta = TM$; equivalently, the distribution is said *bracket-generating*. The fact that by taking iterated brackets of vector fields in the distribution Δ it is possible to obtain all the tangent bundle has very important consequences: every two points in the manifolds can be joined by a horizontal curve. Moreover, as a consequence we can define a distance on M in the usual way:

$$d(p, q) = inf\{\ell(\gamma)\}, \text{where } \gamma(0) = p, \gamma(1) = q,$$

which turns out to be finite since for every couple of points there is a finite-length horizontal curve joining them; this is called *Carathéodory distance*. More precisely, the following theorem holds:

Theorem 1 (Chow-Rashevskii theorem)**.** *Let M be a smooth manifold with a bracket generating distribution Δ. The Carathéodory distance d is finite, continuous and the induced metric topology is the manifold topology.*

2.1.2 The sub-Riemannian tangent space

In order to study local properties of sub-Riemannian structures, we will need their first order approximations, namely the *sub-Riemannian tangent spaces*. In the Riemannian case the first order approximation of the Riemannian manifold (M, g) around $p \in M$ is the differential tangent space T_pM endowed with the scalar product g_p on T_pM; it is a Riemannian manifold itself, and has the structure of a Euclidean space.

Similarly, we expect the sub-Riemannian tangent space to be a sub-Riemannian manifold, which encodes the local data of the sub-Riemannian structure around a point p. This can be realized in many ways: one of them is the *metric tangent space* as explained by Gromov in [15], who gives a procedure to build the tangent space of a general metric space. Roughly, the procedure is the following: given a point p in the metric space (X, d), consider the ball of small radius ε around p $B_\varepsilon(p)$ and apply a dilation with parameter $\frac{1}{\varepsilon}$ to the distance d, in order to let the ball $B_\varepsilon(p)$ be of radius 1 in the new metric. By taking the *Gromov-Hausdorff* limit as $\varepsilon \to 0$ we get the metric tangent space to X at p; it is a metric space that approximates the metric space structure of (X, d) at p. We refer to the paper of Bellaïche [7] for an explanation of the metric tangent space and for more details about the sub-Riemannian tangent space.

We will always assume the regularity of the point p in the following sense:

Definition 3. *A point p on the sub-Riemannian manifold M is **regular** if and only if the growth vector is constant on a neighborhood of p.*

In order to build the nilpotent approximation we need a class of local coordinates that have a good behaviour with respect to the distance. Such coordinates are called *privileged coordinates*, and they always exist on a suitable neighborhood of any regular point $p \in M$. We are not going to give the general definition of privileged coordinates: from now on we will consider 2-step distribution, where a weaker condition turns out to be equivalent.

Definition 4. *Let M be a n-dimensional sub-Riemannian manifold with a rank d distribution Δ, and $p \in M$ a regular point. The local coordinates $(x^1, \dots, x^d, y^1, \dots, y^l)$ are **linearly adapted** at p if $dy^j(\Delta_p) = 0$.*

We will use the shortcut (x, y) for the coordinates $(x^1, \dots, x^d, y^1, \dots, y^l)$ when there will not be any ambiguity. As we already told, the privileged coordinates are always linearly adapted; the converse holds in the 2-step case only. The main property of the privileged coordinates is their behaviour with respect to the distance: indeed we have

$$d(p,q) \asymp |x^1| + \cdots + |x^d| + |y^1|^{1/2} + \cdots + |y^l|^{1/2},$$

where $(x^1, \ldots, x^d, y^1, \ldots, y^l)$ are the coordinates of the point q and $\asymp$ means that the two functions have the same order. Notice that the order $\frac{1}{2}$ of the distance with respect to the coordinates y^j reflects the fact that near to the point p is more difficult to go in their direction.

In the sequel we will need the definition of *weight* of functions and differential operator.

Definition 5 (Weight)**.** *Given the privileged coordinates (x, y) around p, we define the* weight *w such that*

$$w(x^i) = 1, w(y^j) = 2, w\left(\frac{\partial}{\partial x^\alpha}\right) = -1, w\left(\frac{\partial}{\partial y^i}\right) = -2;$$

on the homogeneous monomials the weight is extended additively, for instance

$$w\left(x^b y^j \frac{\partial^3}{\partial x^i \partial^2 y^m}\right) = w(x^b) + w(y^j) + w\left(\frac{\partial}{\partial x^i}\right) + 2w\left(\frac{\partial}{\partial y^m}\right).$$

Moreover, the weight of a differential operator (or of a function) is defined as $\geq s$ if the lowest weight of the monomials in the Taylor expansion is s.

Let us consider now a local trivialization of Δ around p, given by the local vector fields $X_1, \ldots, X_d$. Since they are first order differential operators and $dy^j(X_i) = 0$ for all i and j, it follows that $w(X_i) \geq -1$. For every X_i we define $\widehat{X}_i$ to be its homogeneous part of degree -1. We have the following:

Proposition 2. *The distribution $\widehat{\Delta}$ spanned by the vector fields $\widehat{X}_i$ is bracket generating and has the same growth vector as Δ.*

This new distribution $\widehat{\Delta}$, which is locally a good approximation of the distribution Δ induces also a new distance $\hat{d}$. From the choice of the privileged coordinates we can define the 1-parameter family of dilations centered at p needed to get the metric tangent space:

$$\delta_\lambda(x, y) = (\lambda x, \lambda^2 y).$$

These dilations are defined in order to have the fields $\widehat{X}_i$ and the distance $\hat{d}$ being homogeneous with respect to them:

$$(\delta_\lambda)_* \widehat{X}_i = \lambda^{-1}\widehat{X}_i,\ \hat{d}(p,\delta_\lambda(q)) = \lambda\hat{d}(p,q).$$

Now if we apply the dilations of parameter λ and we rescale the resulting objects by λ^{-1} we obtain:

$$\lambda\,(\delta_\lambda)_* X_i = \widehat{X}_i + \lambda^{-1}R_i(\lambda),$$

where $R_i(\lambda)$ are vector fields of weight ≥ 0 and

$$\lambda^{-1}d(p,\delta_\lambda(q)) = \hat{d}(p,q) + \lambda^{-1}r(\lambda,q),$$

for a suitable function r. This roughly shows that the limit object of X_i and d are $\widehat{X}_i$ and $\hat{d}$ respectively. We found the metric tangent space to be a copy of $\mathbb{R}^n$ with the distance $\hat{d}$; the procedure carries also informations about the structure of Δ, namely the vector fields $\widehat{X}_i$. Such fields together with their brackets give $\mathbb{R}^n$ the structure of a nilpotent Lie algebra $\mathfrak{g}$; moreover, Lie theorem ensures the existence of a simply connected Lie group G with algebra $\mathfrak{g}$. The group G is nothing but a copy of $\mathbb{R}^n$ and in particular is the metric space $(\mathbb{R}^n, \hat{d})$:

Definition 6. *The group G endowed with the left-invariant distribution $\widehat{\Delta}$ and the distance $\hat{d}$ is the* ***sub-Riemannian tangent space*** *of M at p.*

We presented this construction of the sub-Riemannian tangent space because it will be useful in the following; by now we will begin the study starting from a simpler and more algebraic way to recover G. We consider again a regular point $p \in M$ and a set of local vector fields $\xi_1, \ldots, \xi_d$ spanning Δ. We then complete it with a set of vector fields $\psi_1, \ldots, \psi_l$ such that together they give a local frame for the tangent bundle. Then we define the linear space

$$\mathfrak{g} \doteq \Delta_p \oplus T_pM/\Delta_p,$$

together with the basis $X_i \doteq \xi_i(p)$ in Δ_p and $Y_k \doteq \psi_k(p) + \Delta$ in T_pM/Δ_p; then we give a Lie algebra structure to $\mathfrak{g}$ by letting

$$[X_i, X_j] = [\xi_i, \xi_j]_p/\Delta_p = \sum_{k=1}^{l} a_{ij}^k Y_k;$$

moreover Δ_p is endowed with the sub-Riemannian scalar product g_p. We obtained in this way the same nilpotent Lie algebra as before, from which we get the associated Lie group G together with the left-invariant distribution $(L_q)_* \Delta_p$ and the sub-Riemannian metric $(L_q)^* g_p$, where $q \in G$ and L_q is the left multiplication by q. This class of nilpotent Lie groups, which are the local model for the sub-Riemannian manifolds, are also called *Carnot groups*.

2.2 Morse-Bott theory

2.2.1 Morse-Bott functions

We saw in the introduction that the geodesics we study are not isolated critical points of the Energy: this was the case of the Heisenberg group, where the critical set was the disjoint union of circles S^1. In the sequel we will find that this picture is typical: more generally, in a setting where the structure (for instance Riemannian or sub-Riemannian structures) is endowed with families of simmetries and we consider points which are invariant under such simmetries, it is the case that the interesting functionals are invariant under the same simmetries. This forces the critical points to collect themselves into submanifolds; for example on the 2-dimensional sphere, due to the rotational simmetry, the geodesics joining two antipodal points are not isolated but they can be parametrized by circles S^1 in the space of curves. Even though the Morse functions are dense, meaning that after a small perturbation a degenerate function can be approximated by a Morse function, this is not always the best approach because one may lose the simmetries which carry nice information about the structure. Then it is better to consider a more general kind of non-degenerate functions, namely *Morse-Bott functions*. In this section we give a short review of Morse-Bott theory: the interested reader is referred to the original paper by Bott [11] and to the books [17], [24] and [13] for more details (especially for the infinite dimensional case).

Definition 7. *Given a Hilbert manifold X, a smooth function $f : X \to \mathbb{R}$ is a* ***Morse-Bott function*** *if:*

(A) the critical set is the disjoint union of compact smooth manifolds, called ***critical manifolds****;*

(B) if x is a critical point belonging to the critical manifold C then

$$\ker He_x f = T_x C;$$

(C) for every sequence $\{x_k\}$ of critical points on X such that $f(x_k)$ is bounded and $\|\nabla f_{x_k}\| \to 0$, then the sequence $\{x_k\}$ has limit points and every limit point is critical for f.

Condition (C) is usually referred as **Palais-Smale condition** and is automatically satisfied in the finite dimensional case. The smooth manifolds of critical points are called **nondegenerate critical manifolds**; notice that admitting only zero-dimensional critical manifolds we get classical Morse functions.

The second condition is equivalent to the non-degeneracy of the Hessian on the normal space N_xC for $x \in C$. The **index** of the critical manifold C is defined as the maximum of the dimensions of subspaces $V \subset N_xC$ where the Hessian is negative-definite; since the Hessian is nondegenerate in the all normal bundle, this number does not depend on the point $x \in C$ and it is denoted by $\mathrm{ind}(C)$.
Under this assumptions it is still possible to describe what happens to the topology of the sublevels $X^c \doteq f^{-1}(-\infty, c)$ of the Morse-Bott function f by increasing c.

Let us consider a Morse-Bott function $f : X \to \mathbb{R}$: the first fundamental theorem of Morse theory describe how the sublevels (don't) change when c increases without passing critical values:

Theorem 3. *Let $f : X \to \mathbb{R}$ be a Morse-Bott function on X. If the interval $[a, b] \subset \mathbb{R}$ does not contain critical values, X^b is diffeomorphic to X^a.*

This can be proved by introducing a *gradient-like vector field*, namely a vector field V such that $\langle f, V\rangle > 0$; then X^b is deformed to X^a along the integral curves of the flow of V. Moreover if we let a be a critical value for the Morse-Bott function f the sublevels aren't diffeomorphic one to another anymore, but still there exists a deformation.

Theorem 4. *If $f : X \to \mathbb{R}$ is a C^1 function satisfying the Palais-Smale condition and $(a, b] \subset \mathbb{R}$ does not contain critical values, X^a is a strong deformation retract of X^b.*

For the proof see Lemma 3.2 in Chapter 1 of [13].

It only remains to recall what happens when we pass a critical value for a Morse-Bott function. Given a critical manifold C we restrict the tangent bundle TX to C and consider a sub-bundle $E^-_C \subset TC$ such that the Hessian of f is negative definite and the fibers have dimension equal to $\mathrm{ind}(C)$, together with the unit disk bundle $D^-_C \subset E^-_C$. With this notation the following theorem generalizes the classical one (the statement we present here is actually the one in [17]).

Theorem 5 (Bott)**.** *Let $f : X \to \mathbb{R}$ be a Morse-Bott function, and let c be a critical value. For $\delta > 0$ sufficiently small the sublevel $X^{c+\delta} = X \cap \{f \leq c + \epsilon\}$ is homotopic to the sublevel $X^{c-\epsilon}$ with the unit disk bundle D^-_C glued along the boundary.*

Since in our setting we consider homology with $\mathbb{Z}_2$ coefficients, it follows from the Thom isomorphism that:

$$H_*(X^{c+\delta}, X^{c-\delta}) \simeq H_*(D^-_C, \partial D^-_C),$$

where in the last equation we allow the critical manifold C to be non-connected, in which case we actually have a disjoint union of different bundles (with possibly different rank, corresponding to the possibly different indexes of the components of C).

2.2.2 Morse-Bott inequalities

One of the main tools in Morse theory are the *Morse inequalities*: they allow to compare the number of critical points to the topological properties of the manifold where a Morse function is given. We denote with $b_k(X)$ the k-th Betti number, namely the number of independent cycles of dimension k in the homology of the topological space X.

Proposition 6 (Morse inequalities)**.** *Let X be a smooth (Hilbert) manifold and $f : X \to \mathbb{R}$ be a Morse function:*

$$b_k(X) \leq \#\{\textit{critical points with index } k\}.$$

Alternatively we can state *Morse inequalities* in terms of the Poincaré polynomial of X and Morse polynomial of f; we recall the former to be

$$P_X(t) = \sum_k b_k t^k,$$

while the latter is defined as

$$M_f(t) = \sum_x t^{\mathrm{ind}(x)},$$

where x runs through the critical points of f. Moreover, we define the following partial order relation in the set of polynomials of one variable:

$$P(t) \prec Q(t)$$

if there exists a polynomial $R(t)$ with non-negative coefficients such that

$$Q(t) = P(t) + (1+t)R(t).$$

we are ready to state the following: The Morse inequalities can now be written as follows:

$$P_X(t) \prec M_f(t).$$

It is easy to see how the classical Morse inequalities follow from the polynomial ones; even though the inequalities written in terms of polynomials may appear fancy, they are quite useful to state the Morse inequalities in the Morse-Bott case.

In the Morse-Bott setting the *Morse-Bott polynomial* is defined in terms of Poincaré polynomials of the non-degenerate critical manifolds:

$$M_f(t) = \sum_C P_C(t)\, t^{\mathrm{ind}(C)},$$

where C runs through the critical manifolds. Now, the Morse-Bott inequalities have the same expression as the Morse inequalities, namely

Proposition 7 (Morse-Bott inequalities)**.** *Given a Morse-Bott function f on the Hilbert manifold X, we have*

$$P_X(t) \prec M_f(t).$$

What we will need in the following is a more rough inequality which doesn't depend on the indexes; by putting in the polynomials $t = 1$, from the Morse-Bott inequalities follows the relation

$$b(X) \leq \sum_C b(C),$$

where by $b(M)$ we denote the *total Betti number* of the topological space M, namely the sum of all the Betti numbers.

Just a final remark before going to the next section: this theory keeps true with obvious modifications to the notation if we consider the sublevels $X^c = \{f \leq c\}$ of the Morse-Bott function f instead of the whole manifold.

2.3 Semialgebraic geometry

2.3.1 Semi-algebraic sets

In this subsection we will give a couple of basic definitions from semialgebraic Geometry, and we will state some properties that will be useful in the following. For a detailed exposition of semialgebraic geometry we refer to [8] and [10].

To begin with, we give the definition of algebraic and semialgebraic sets. Real algebraic sets are defined in the same way as complex ones:

Definition 8 (Algebraic set)**.** *Given a subset $B \subset \mathbb{R}[x_1, \ldots, x_n]$, the zero locus*

$$S \doteq \{x \in \mathbb{R}^n \mid \exists f \in B \ s.t.\ f(x) = 0\}$$

is called ***algebraic set****.*

Since $\mathbb{R}$ is an ordered field, it is possible to define subsets of $\mathbb{R}^n$ with inequalities as well:

Definition 9 (Semialgebraic set)**.** *A **semialgebraic subset** S of $\mathbb{R}^n$ is a subset of the form*

$$\bigcup_{i=1}^{s}\bigcap_{j=1}^{r_i}\{x \in \mathbb{R}^n \mid f_{ij} \bowtie_{ij} 0\},$$

where $f_{ij} \in \mathbb{R}[x_1, \ldots, x_n]$ and $\bowtie_{ij}$ may be the binary relations $=$ or $<$.

It turns out that only a finite number of polynomials are required in order to define a semialgebraic set.

Proposition 8. *Every semialgebraic set $S \in \mathbb{R}^n$ can be written as the finite union of subsets of the form*

$$\{x \in \mathbb{R}^n \mid f_1(x) = 0, \ldots, f_l(x) = 0, g_1(x) < 0, \ldots, g_m(x) < 0\}$$

where $f_1, \ldots, f_l, g_1, \ldots, g_m \in \mathbb{R}[x_1, \ldots, x_n]$.

Semialgebraic sets behave well with respect to topological operations. Indeed, we have

Proposition 9. *The closure and the interior of a semialgebraic set are still semialgebraic sets.*

Since semialgebraic sets are defined by real polynomials, they may not be smooth; though, they can be split into the disjoint union of smooth manifolds:

Definition 10 (Stratification)**.** *Given a semialgebraic set S in $\mathbb{R}^n$, a **stratification** of S is a finite partition $S = \bigsqcup S_j$ such that:*

- *every S_j is a smooth submanifold of S;*
- *the closure of every S_j is the union of S_j together with other subsets of the same partition of lower dimension.*

It is important to recall that every semialgebraic set admits a stratification; the elements of the partition giving a stratification are called *strata*. Semialgebraic sets are provided with their own class of maps, namely:

Definition 11 (Semialgebraic maps)**.** *A map $f : S_1 \to S_2$ between two semialgebraic sets $S_1 \subset \mathbb{R}^n$ and $S_2 \subset \mathbb{R}^m$ is **semialgebraic** if its graph in $\mathbb{R}^{m+n}$ is a semialgebraic set.*

A remarkable property of semialgebraic sets is that they are closed with respect to the projections. Indeed, choose a copy of $\mathbb{R}^n$ in $\mathbb{R}^{n+1}$ and let p be the projection $p : \mathbb{R}^{n+1} \to \mathbb{R}^n$:

Theorem 10. *The image of any semialgebraic set $S \subset \mathbb{R}^{n+1}$ under the projection $p : \mathbb{R}^{n+1} \to \mathbb{R}^n$ is still a semialgebraic set.*

We will need in the following an improvement of the Sard's Lemma in the realm of semialgebraic sets:

Theorem 11 (Semialgebraic Sard's Lemma)**.** *Let $f : S_1 \to S_2$ be a semialgebraic smooth map between two semialgebraic smooth manifolds S_1 and S_2. The set of critical values of f is a semialgebraic subset of S_2 with dimension strictly smaller that the dimension of S_2.*

We conclude this subsection with another important property of semialgebraic sets, which allows to reduce the preimage of a semialgebraic set to a finite union of cartesian products.

Theorem 12 (Hardt's triviality)**.** *Let $S_1 \subset \mathbb{R}^n$ and $S_2 \subset \mathbb{R}^m$ be semialgebraic sets, and $f : S_1 \to S_2$ a continuous semialgebraic map. There exists a finite partition T_i of T such that for every $x_i \in T_i$ the preimage $f^{-1}(T_i)$ is semialgebraically homeomorphic to the product $T_i \times f^{-1}(x_i)$.*

2.3.2 Cohomology of the intersection of real quadrics

We will see in the following that the study of the horizontal path spaces will be reduced to the study of intersection of real quadrics in finite dimension by homotopy equivalence; thus in this section we present useful results from [5, 18, 19] for the study of the Betti numbers of the intersection of real quadrics on the sphere S^n. The motivating example is the case of the zero locus Y of one single nondegenerate quadratic form q on the sphere S^n: if $\mathrm{i}^-(q)$ denotes the negative inertia index of q, then:

$$Y \simeq S^{\mathrm{i}^-(q)-1} \times S^{n-\mathrm{i}^-(q)}.$$

In particular we see that the knowledge of the *index* function on the whole line spanned by q in the space of all quadratic forms determines the topology (since by non-degeneracy $n - \mathrm{i}^-(q) = \mathrm{i}^-(-q) - 1$).
More generally if we have l quadratic forms $q_1, \ldots, q_l$ in $n+1$ variables, we need to study the index function on the linear span of the quadratic forms

$q_1, \ldots, q_l$ in the space of all quadratic forms $\mathrm{Sym}(n)$. To this purpose we define the map

$$\begin{array}{rcl} \mathbb{R}^l & \to & \mathrm{Sym}(n) \\ \eta = (\eta_1, \ldots, \eta_l) & \mapsto & \eta q \doteq \eta_1 q_1 + \cdots + \eta_l q_l, \end{array} \tag{2.1}$$

where $q = (q_1, \ldots, q_l)$ is the vector whose entries are the given quadratic forms. Then we consider the function $\eta \mapsto \mathrm{i}^-(\eta q)$; in the generic case this function is the restriction of the negative inertia index function to the span of $q_1, \ldots, q_l$ in the space of all quadratic forms. Although the general theory is more detailed, for our purposes we need explicit computations only in the case $l = 2$ and in the general case it will suffice to have quantitative bounds on the topology of:

$$Y = \{x \in S^n \,|\, q_1(x) = \cdots = q_l(x) = 0\}.$$

In the case $l = 2$ we consider a unit circle S^1 in $\mathbb{R}^2$ and the restriction $\mathrm{i}^-|_{S^1}$; also for $j \geq 0$ we let:

$$P_j = \{\eta \in S^1 \,|\, \mathrm{i}^-(\eta q) \leq j\}.$$

The following formula (2.2) is proved in [18] and relates the Betti numbers of Y to the topology of the sets P_j; we denote by $\tilde{b}_j(Y)$ the rank of $\tilde{H}^j(Y; \mathbb{Z}_2)$. The general bound (2.3) for the topology of Y in the case $l \geq 2$ is proved in [19]. The reader is referred to [5, 18, 19] for more details.

Proposition 13. *If Y is the intersection of* two *quadrics on the sphere S^n and $0 \leq j \leq n-3$, then:*

$$\tilde{b}_j(Y) = \tilde{b}_{n-j-1}(S^n \backslash Y) = b_0(P_{j+1}, P_j) + b_1(P_{j+2}, P_{j+1}). \tag{2.2}$$

Moreover if Y is defined by $l \geq 2$ quadratic equations on S^n, on $\mathbb{R}P^n$ or in $\mathbb{R}^n$, then:

$$b(Y) \leq O(n)^{l-1}. \tag{2.3}$$

It is possible to apply the above technique also in the case Y is the intersection of quadrics on the unit sphere in some (possibly infinite dimensional) Hilbert space H. The main differences for this infinite dimensional case are the following: Y must be non-singular; $\check{b}_i$ denotes the rank of the i-th *Cech* cohomology group; the negative inertia index might be infinite for some $\eta \in S^1$, but these η are already excluded by the condition $\mathrm{i}^-(\eta) \leq j < \infty$. With these modification we have the following result from [1]; formula (2.5) is the analogue of (2.2), but the condition that H is infinite dimensional allows to remove the restriction on the range for j.

Theorem 14. *Let* $q_1, \ldots, q_l$ *be* continuous *quadratic forms on the Hilbert space* H *and* Y *be their (nondegenerate) common zero locus on the unit sphere. Then:*

$$H_*(Y) = \varinjlim_{V \in \mathcal{F}} \{H_*(Y \cap V)\}. \tag{2.4}$$

where $\mathcal{F}$ *denotes the family of all* finite *dimensional subspaces of* H*. Moreover in the case* $l = 2$*:*

$$\tilde{b}_j(Y) = \check{b}_0(P_{j+1}, P_j) + \check{b}_1(P_{j+2}, P_{j+1}) \tag{2.5}$$

Sketch. Since Y is assumed to be non-singular, then it has a tubular neighborhood U in H and $H_*(Y) \simeq H_*(U)$. In particular every singular chain in U is homotopic to one whose image is contained in a finite dimensional subspace and (2.4) follows. To prove (2.5) we fix a $j \geq 0$; then using (2.2) we have:

$$\tilde{b}_j(Y \cap V) = b_0(P_{j+1}(V), P_j(V)) + b_1(P_{j+2}(V), P_{j+1}(V))$$

where $V \subset H$ is a sufficiently big finite dimensional subspace (the condition required on the dimension is $\dim(V) - 3 \geq j$) and $P_j(V) = \{\eta \in S^1 \,|\, \mathrm{i}^-(\eta q|_V) \leq j\}$. Now the sets $\{P_j(V)\}_{V \in \mathcal{F}}$ are also partially ordered by inclusion: if $V_1 \subset V_2$, then $P_j(V_2) \subset P_j(V_1)$. It is not difficult to show that under the isomorphism

$$\tilde{H}_j(Y \cap V) \simeq H^0(P_{j+1}(V), P_j(V)) \oplus H^1(P_{j+2}(V), P_{j+1}(V))$$

the inclusion morphism on the homology $H_j(Y \cap V_1) \to H_j(Y \cap V_2)$ is induced by the restriction morphism (see [1]):

$$\bigoplus_{i=0,1} H^i(P_{j+i+1}(V_1), P_{j+i}(V_1)) \to \bigoplus_{i=0,1} H^i(P_{j+i+1}(V_2), P_{j+i}(V_2))$$

Since the sets $\{P_j(V)\}_{V \in \mathcal{F}}$ are Euclidean Neighborhood Retracts (being semialgebraic sets), then by the continuity property of Cech cohomology:

$$\tilde{H}_*(Y) = \varprojlim_{V \in \mathcal{F}} \{H^*(P_{j+1}(V), P_j(V))\} = \check{H}^*\big(\bigcap_{V \in \mathcal{F}} P_{j+1}(V), \bigcap_{V \in \mathcal{F}} P_j(V)\big).$$

Finally P_j equals by construction $\bigcap_{V \in \mathcal{F}} P_j(V)$ and the conclusion follows. $\square$

2.3.3 A useful stratification of $\mathfrak{so}(d)$

As it will be explained in the next chapter, the possible structures of rank d Carnot groups are totally encoded in the linear subspaces of $\mathfrak{so}(d)$, the space of skew-symmetric matrices on $\mathbb{R}^d$. In the following chapters, transversality arguments will be important tools in order to prove our results; therefore we construct a stratification of $\mathfrak{so}(d)$ and we compute the codimensions of such strata, generalizing the results from the appendix of [9].

We are interested in studying the dimensions of the semi-algebraic sets with generalized eigenvalues of given multiplicities and with given dimension of the kernel. Every skew-symmetric matrix A can be written in its canonical form as a block matrix, with blocks on the diagonal of the form

$$\alpha J_2 = \begin{pmatrix} 0 & \alpha \\ -\alpha & 0 \end{pmatrix},$$

(J_2 being the canonical symplectic matrix in $\mathfrak{so}(2)$) and a 0-block of the dimension of the kernel. By *generalized eigenvalues* we mean the entries like α.

We introduce the set:

$$\Gamma_{k|m_1,\dots,m_r} \subset \mathfrak{so}(d)$$

defined to be the set of skew-symmetric matrices in $\mathfrak{so}(d)$ with: (a) dimension of the kernel equal to k and (b) multiplicities of the generalized eigenvalues $m_1, \dots, m_r$ (with $m_1 \geq m_2 \geq \cdots m_r$).

By acting with $SO(d)$ on a matrix $A \in \Gamma_{k|m_1,\dots,m_r}$ with eigenvalues $\alpha_1, \dots, \alpha_r$ (corresponding to the ordered multiplicities), we can put it in the form:

$$\mathrm{Diag}\,(\,\alpha_1 J_{2m_1}, \dots, \alpha_r J_{2m_r}, 0_k)\ ,$$

Let us look at the stabilizer $SO(d)_A$ of A: first of all it has to fix every eigenspace of A, since eigenspaces with different eigenvalues are orthogonal. On the kernel K the stabilizer is the restriction of $SO(d)$ on K, i.e. a copy of $SO(k)$; on the eigenspace with eigenvalue α_i the restriction of $SO(d)$ is a copy of $SO(2m_i)$, but the stabilizer has to fix the symplectic matrix; it follows that the stabilizer act as a copy of:

$$SO(2m_i) \cap Sp(2m_i) = U(m_i).$$

where J_{2n} is the symplectic matrix in $\mathbb{R}^{2n}$ and 0_k is the null matrix on $\mathbb{R}^k$.

Now it is possible to compute the codimension of the orbit $\mathrm{Ad}(SO(d))A$ of A by the adjoint action Ad of $SO(d)$ on $\mathfrak{so}(d)$ with known stabilizer:

$$\begin{aligned} \operatorname{codim}\mathrm{Ad}(SO(d))A &= \dim\mathfrak{so}(d) - \dim\mathrm{Ad}(SO(d))A = \\ &= \dim\mathfrak{so}(d) - \dim SO(d) + \dim SO(d)_A = \dim SO(d)_A\,. \end{aligned}$$

We know that the stabilizer $SO(d)_A$ is:

$$SO(d)_A = SO(k) \times U(m_1) \times \ldots \times U(m_r)\,,$$

and its dimension is:

$$\dim SO(d)_A = \frac{k(k-1)}{2} + \sum_{i=1}^{r} m_i^2.$$

Let us now consider the eigenvalues α_i: as long as they are distinct (so they preserve their multiplicities) they are smooth functions of the matrices [16]. On the set $\Gamma_{k|m_1,\ldots,m_r}$ this condition holds true (by definition), hence we have a smooth map:

$$\psi : \Gamma_{k|m_1,\ldots,m_r} \to \mathbb{R}^r$$

given by $A \mapsto (\alpha_i)$. This map is indeed a submersion on the open subset $\mathcal{O}$ of vectors in $\mathbb{R}^r$ with distinct entries. The fibers of the map ψ are the orbits of the adjoint action and they are diffeomorphic to a fixed manifold $SO(d)/SO(d)_A$; in particular $\Gamma_{k|m_1,\ldots,m_r}$ is a fiber bundle over $\mathcal{O}$ with fibers diffeomorphic to $SO(d)/SO(d)_A$. Now we can compute the codimension of $\Gamma_{k|m_1,\ldots,m_r}$ in $\mathfrak{so}(d)$:

$$\operatorname{codim}_{\mathfrak{so}(d)}\Gamma_{k|m_1,\ldots,m_r} = \operatorname{codim}_{\mathfrak{so}(d)}\mathrm{Ad}(SO(d))A - r = \frac{k(k-1)}{2} - r + \sum_{i=1}^{r} m_i^2 \tag{2.6}$$

$$= \frac{k(k-1)}{2} + \sum_{i=1}^{r}\left(m_i^2 - 1\right). \tag{2.7}$$

Since we are interested in the matrices with integer eigenvalues, we will need to stratify $\Gamma_{k|m_1,\ldots,m_r}$ in infinite semialgebraic sets with given integer eigenvalues.

Later we will need to deal with sets of skew-symmetric matrices with at least one imaginary integer eigenvalues; we fix a stratum $\Gamma_{k|m_1,\ldots,m_r}$ and we proceed as follows.

Given $\vec{n} = (n_1, \ldots, n_r) \in \mathbb{N}^r$ with non-negative entries such that all the non-zero entries are distinct, by $\Gamma_{k|m_1,\ldots,m_r|\vec{n}}$ we will mean the subset of $\Gamma_{k|m_1,\ldots,m_r}$ with the eigenvalue of multiplicity m_j equal to in_j if and only if $n_j > 0$. The eigenvalues corresponding to zero entries of $\vec{n}$ vary in $\mathbb{R}$. Since by fixing an eigenvalue we drop the dimension of the subset by 1, we have the following:

Proposition 15. *Given $\vec{n} \in \mathbb{N}^r$ with the properties described above and with ν non-zero entries, the submanifold $\Gamma_{k|m_1,\dots,m_r|\vec{n}}$ has codimension ν in $\Gamma_{k|m_1,\dots,m_r}$, thus its codimension in the set of all matrices $\mathfrak{so}(d)$ is*

$$\operatorname{codim}_{\mathfrak{so}(d)}\Gamma_{k|m_1,\dots,m_r|\vec{n}} = \frac{k(k-1)}{2} + \sum_{j=1}^{r}\left(\mu_j^2 - 1\right) + \nu. \tag{2.8}$$

Chapter 3

Horizontal path spaces

3.1 Geodesics in step 2 Carnot groups

3.1.1 Step 2 Carnot groups

Now we will focus mainly on Carnot groups, not just as sub-Riemannian tangent spaces but as sub-Riemannian manifolds in their own right.

In order to fix the notation, we recall that a *step 2 Carnot group* is a connected, simply connected Lie group G whose Lie algebra $\mathfrak{g} = T_eG$ decomposes as:

$$\mathfrak{g} = \Delta \oplus \Delta^2, \quad \text{with} \quad [\Delta, \Delta] = \Delta^2, \quad [\Delta, \Delta^2] = 0 \quad \text{and} \quad [\Delta^2, \Delta^2] = 0;$$

as vector spaces here we have $\Delta \cong \mathbb{R}^d$ and $\Delta^2 \cong \mathbb{R}^l$. We recall that under the above assumption on the structure of $\mathfrak{g}$, the exponential map $\exp : \mathfrak{g} \to G$ is an analytic diffeomorphism, hence in particular $G \simeq \mathbb{R}^{d+l}$.

The geometric structure on G is given by fixing a scalar product h on Δ and considering the distribution $\Delta_q = (L_q)_*\Delta$ together with the extension of h by left translation.

Two such Carnot groups (G_1, Δ_1, h_1) and (G_2, Δ_2, h_2) are considered to be isomorphic if there exists a Lie algebra isomorphism $L : \mathfrak{g}_1 \to \mathfrak{g}_2$ such that $L\Delta_1 = \Delta_2$ and $L|^*_{\Delta_1} h_2 = h_1$: in fact by the simple connectedness assumption the linear map L integrates to a Lie group isomorphism $\phi : G_1 \to G_2$ and the global geometric structures are related by $(G_1, \Delta_1, h_1) = (G_1, \phi^{-1}\Delta_2, \phi^* h_2)$.

Now we fix an *orthonormal* basis $\{e_1, \ldots, e_d\}$ of Δ and a basis $\{f_1, \ldots, f_l\}$ of Δ^2, so that the bracket structure can be written as:

$$[e_i, e_j] = \sum_{k=1}^{l} a_{ij}^k f_k, \quad \text{for all } i, j \in \{1, \ldots, d\} \tag{3.1}$$

where each matrix $A_k = (a_{ij}^k)$ belongs to $\mathfrak{so}(d)$; from these matrices we get the vector space

$$W = \operatorname{span}\{A_1, \ldots, A_l\} \subset \mathfrak{so}(d).$$

Notice that the step 2 condition implies that the matrices A_i are linearly independent in $\mathfrak{so}(d)$: otherwise we would have a null linear combination $\lambda_1 A_1 + \cdots \lambda_l A_l = 0$ where at least one of the λ_i is non zero. This would mean that for every couple of vector $X = \sum_{i=1}^d X^i e_i, Y = \sum_{j=1}^d Y^j e_j$ in Δ, their bracket would be annihilated by the covector $\lambda \doteq (\lambda_1, \ldots, \lambda_l) \in (\Delta^2)^*$:

$$\begin{aligned}\langle \lambda, [X,Y] \rangle = \sum_{i,j=1}^{d} X^i Y^j \langle \lambda, [e_i, e_j] \rangle = \sum_{i,j} \sum_{k=1}^{l} X^i Y^j a_{ij}^k \langle \lambda, f_k \rangle = \\ = \sum_{i,j=1}^{d} \sum_{k=1}^{l} X^i Y^j \lambda_k a_{ij}^k = \sum_{k=1}^{l} \lambda_k A_k(X,Y) = 0,\end{aligned}$$

and so $[\Delta, \Delta] \neq \Delta^2$, contradicting the step 2 hypotesis.

So we can invert the point of view and see how each vector space W of dimension l in $\mathfrak{so}(d)$ gives $\mathbb{R}^{d+l}$ a step 2 Carnot group structure. Given a basis $e_1, \ldots, e_d, f_1, \ldots, f_l$ on $\mathbb{R}^{d+l}$ and a basis $\{A_1, \ldots, A_l\}$ for W we define the Lie brackets of the algebra as

$$[e_i, e_j] = \sum_{k=1}^{l} A_k(e_i, e_j) f_k = \sum_{k=1}^{l} a_{ij}^k f_k.$$

Moreover we define the sub-Riemannian scalar product by declaring $e_1, \ldots, e_d$ to be orthonormal.

Proposition 16. *The isomorphism class of the Carnot group does not depend on the choice of the basis of W.*

Proof. Let $\{A'_1, \ldots, A'_l\}$ be another basis of W and $M = (m_h^k)$ the basis-change matrix, such that $A_k = \sum_{k=1}^l m_h^k A'_h$. Now we can build another Carnot group by defining its Lie algebra $\mathfrak{g}'$ on $\mathbb{R}^{d+l}$ by finding a basis $\{e'_1, \ldots, e'_d, f'_1, \ldots, f'_l\}$ and setting $\{e'_1, \ldots, e'_d\}$ to be orthonormal; the structure constants are given by the entries of the matrices A'_h as in equations (3.1). The map ϕ, defined on the basis elements by $e_i \mapsto e'_i$ and $f_k \mapsto \sum_{k=1}^l m_h^k f_k = f'_h$, gives an isomorphism between $\mathfrak{g}'$ and $\mathfrak{g}$: it preserves the scalar product and the Lie algebra brackets, since

$$
\begin{aligned}
[\phi(e_i), \phi(e_j)] = [e_i', e_j'] &= \sum_{h=1}^{l} a'^{h}_{ij} f_h' = \sum_{h,k=1}^{l} a^k_{ij} m^h_k f_h' \\
&= \sum_{k=1}^{l} a^k_{ij} \phi(f_k) = \phi\left([e_i, e_j]\right).
\end{aligned}
$$

□

Remark 1 (The moduli space of Carnot Groups). Once a left-invariant sub-Riemannian structure is given, changing $\{e_1, \ldots, e_d\}$ to another orthonormal basis $\{Me_1, \ldots, Me_d\}$ (where M is an orthogonal matrix in $O(d)$) changes W into $W' = MWM^T$. Thus denoting by $G(l, \mathfrak{so}(d))$ the Grassmannian of l-planes in $\mathfrak{so}(d)$, the (naive) moduli space of step two Carnot Groups is represented by the quotient $\mathcal{M}_{l,d} = G(l, \mathfrak{so}(d))/O(d)$. Since $\mathcal{M}_{l,d}$ is the quotient of a manifold by a Lie group action, the quotient map is open and perturbing W defines a "genuine" perturbation of the isomorphism class of the corresponding Carnot group; in particular this means that a generic choice of W results in a generic choice of an isomorphism class of Carnot groups.

Motivated by the above remark, once $W \in G(l, \mathfrak{so}(d))$ is fixed we consider the Carnot group given by exponentiating $\mathfrak{g} = \mathbb{R}^d \oplus \mathbb{R}^l$, whose Lie algebra is given as follows: $\{e_1, \ldots, e_d\}$ is the standard orthonormal basis for $\mathbb{R}^d$, $\{f_1, \ldots, f_l\}$ is the standard basis for $\mathbb{R}^l$ and fixing a basis $\{A_1, \ldots, A_l\}$ for W, the Lie brackets are given by equation (3.1); we will call W the *Carnot group structure*.

The following theorem gives a geometric realization of Carnot groups.

Theorem 17. *Let $\{A_1, \ldots, A_l\}$ be a basis for $W \subset \mathfrak{so}(d)$ and for $i = 1, \ldots, d$ consider the vector fields E_i on $\mathbb{R}^{d+l}$ defined in coordinates (x, y) by:*

$$
E_i(x, y) = \frac{\partial}{\partial x_i}(x, y) - \frac{1}{2} \sum_{k=1}^{l} \sum_{j=1}^{d} a^k_{ij} x_j \frac{\partial}{\partial y_k}(x, y).
$$

Then the sub-Riemannian manifold $(\mathbb{R}^{d+l}, \Delta = span\{E_1, \ldots, E_d\}, g)$, where g is the standard Euclidean metric, is isomorphic to the Carnot group defined by W.

Notice that these vector fields are homogeneous of weight -1 (the weight given in the definition 5) and the coordinates (x, y) are automatically privileged since $dy^k(E_i) = -\frac{1}{2}\sum_{j=1}^{d} a^k_{ij} x_j$ which is equal to 0 in the origin.

3.1.2 The End-point map

From now on we will consider horizontal paths parametrized on the interval $I = [0, 2\pi]$ in order to simplify notation: we will later need to expand the components of an horizontal path into their Fourier series, and a different choice of the interval will produce a completely equivalent theory.

As already told, the bracket generating condition implies that any two points in G can be joined by an horizontal path: we define Ω to be the set of *all* the horizontal paths starting from the identity e of G. The set Ω can be given a *Hilbert manifold structure* as follows. Let $u = (u_1, \ldots, u_d) \in L^2(I, \mathbb{R}^d)$ and consider the Cauchy problem:

$$\dot{\gamma}(t) = \sum_{i=1}^{d} u_i(t) E_i(\gamma(t)), \quad \gamma(0) = e.$$

This is just a ODE problem set on $\mathbb{R}^{d+l}$, using the trivialization given by the global vector fields E_i; in this case the identity element $e \in G$ corresponds to the zero of $\mathbb{R}^{d+l}$. By Carathéodory's Theorem the above Cauchy problem has a local solution γ_u and we consider the set:

$$\mathcal{U} = \{u \in L^2(I, \mathbb{R}^d) \,|\, \gamma_u \text{ is defined for } t = 2\pi\}.$$

For general sub-Riemannian manifolds, $\mathcal{U}$ is an open subset of $L^2(I, \mathbb{R}^d)$ (by ODE's continuous dependence theorem) and is called the set of *controls*; in our case the estimates for the final time can be made uniform (the constraints on the velocities are linear) and we actually have $\mathcal{U} = L^2(I, \mathbb{R}^d)$. Associating to each u the corresponding path γ_u gives thus a global coordinate chart and by slightly abusing of notation in the sequel we will often identify Ω with $\mathcal{U}$.

Once we are given the Carnot group structure $W = \operatorname{span}\{A_1, \ldots, A_l\}$, we can use Theorem 17 to write down the above ODE in a more explicit form:

$$\begin{cases} \dot{x} = u \\ \dot{y}_i = \frac{1}{2} x^T A_i u \end{cases} \quad \text{and} \quad \gamma(0) = 0.$$

In this framework the *End-point map* is the *smooth* map:

$$F : \Omega \longrightarrow G,$$

that associates to each curve γ its final point $\gamma(2\pi)$.

We define the vector of matrices $A = (A_1, \ldots, A_l)$, and we can use again Theorem 17 and write the End-point map as:

$$F(u) = \left(\int_I u(t)\, dt, \frac{1}{2} \int_I \left\langle \int_0^t u(\tau) d\tau, A\, u(t) \right\rangle dt \right); \tag{3.2}$$

(here the brackets denote the sub-Riemannian scalar product on the Lie algebra $\mathfrak{g}$).

In the sequel we will mainly be interested in horizontal paths whose endpoints lie on a particular submanifold of G where we will experience the most typical properties of the sub-Riemannian geometries. Being G of step two, we know that Δ^2 is an abelian subalgebra of $\mathfrak{g}$; therefore we can identify Δ^2 with the submanifold $\exp(\Delta^2) \subset G$ and using Theorem 17 we can write this identification as:

$$\xi_1 f_1 + \ldots + \xi_l f_l \mapsto (\underbrace{0, \ldots, 0}_{x}, \underbrace{\xi_1, \ldots, \xi_l}_{y}),$$

(here as above $\{f_1, \ldots, f_l\}$ is a basis of Δ^2).

Remark 2. By abuse of notation we will denote $\exp(\Delta^2) \subset G$ as Δ^2, since by the above identification they are essentially the same; we are motivated by the fact that we will need covectors in $(\Delta^2)^*$ to act on the points in $\exp(\Delta^2)$ via the previous identification. So in the following we will consider points in Δ^2 instead of points in $\exp(\Delta^2)$.

Definition 12. *The points $p \in \Delta^2$ are called **vertical**.*

We study now the structure of the set of horizontal paths whose endpoints are vertical. It turns out that in the coordinates given by the controls in $\mathcal{U}$ it coincides with the kernel of the differential of F at $0 \in \mathcal{U}$

$$H = \ker D_0 F,$$

as described by the following proposition.

Proposition 18. *The following properties hold:*

(a) $H = \left\{u \in L^2(I, \mathbb{R}^d) \mid \int_I u \, dt = 0\right\}$*;*

(b) $u \in H \Leftrightarrow F(u) \in \Delta^2$*;*

(c) $F|_H = Hess_0 F$*.*

Proof. For point (a) we compute the differential $D_0 F$: by taking a variation εv of the constant curve $\gamma \equiv 0$ we easily see that

$$D_0 F v = \frac{d}{d\varepsilon}\bigg|_{\varepsilon=0} \left(\varepsilon \int_I v(t) \, dt, \frac{1}{2}\varepsilon^2 \int_I \left\langle \int_0^t v(\tau) d\tau, A \, v(t) \right\rangle dt \right) = \\ = \left(\int_I v(t) \, dt, 0 \right) \in \mathfrak{g}.$$

which proves property (a).

Point (b) is a direct consequence of equation (3.2).

For point (c) we notice that the Hessian $\mathrm{He}_0 F$ is defined on $H = \ker D_0F$ with values in $\mathrm{coker} D_0F = \Delta^2$, thus has the same range as $F|_H$. As in the proof of (a) if we consider the second derivative of a variation and we easily obtain $\mathrm{He}_0 F$ has the same expression of F when restricted to H. $\square$

We denote by q the Hessian of F at zero, i. e. the quadratic map:

$$q \doteq F|_H : H \to \mathbb{R}^l.$$

Every component q_i of q is a quadratic form on H and its explicit expression is given by:

$$q_i(u) = \frac{1}{2}\int_I \left\langle \int_0^t u(\tau)d\tau, A_i\, u(t) \right\rangle dt. \tag{3.3}$$

By polarization we obtain the expression for the associated bilinear form:

$$\begin{aligned} q_i(u,v) &= \frac{1}{4}\left(\int_I \left\langle \int_0^t u(\tau) + v(\tau)d\tau, A_i(u(t)+v(t)) \right\rangle dt + \right. \\ &\left. - \int_I \left\langle \int_0^t u(\tau)d\tau, A_i\, u(t) \right\rangle dt - \int_I \left\langle \int_0^t v(\tau)d\tau, A_i\, v(t) \right\rangle dt \right) = \\ &= \frac{1}{4}\left(\int_I \left\langle \int_0^t u(\tau)d\tau, A_i\, v(t) \right\rangle dt + \int_I \left\langle \int_0^t v(\tau)d\tau, A_i\, u(t) \right\rangle dt \right) = \\ &= \frac{1}{4}\left(\int_I \left\langle \int_0^t u(\tau)d\tau, A_i\, v(t)dt \right\rangle - \int_I \left\langle v(t), A_i \int_0^t u(\tau)d\tau \right\rangle dt \right) = \\ &= \frac{1}{2}\int_I \left\langle \int_0^t u(\tau)d\tau, A_i\, v(t) \right\rangle dt, \end{aligned}$$

where the fourth row follows from integration by parts.

Moreover to every q_i there corresponds a symmetric operator $Q_i : H \to H$ defined by:

$$q_i(u) = \langle u, Q_i u\rangle_H \quad \text{for all } u \in H.$$

(Recall that $\langle u,v\rangle_H = \int_I \langle u,v\rangle dt$). We will use the notation Q for the map $(Q_1,\ldots,Q_l) : H \to H \otimes \mathbb{R}^l$.

Remark 3 (Composition of vector valued maps with covectors). In the following we will extensively make use of the following identifications. Given $A = (A_1,\ldots,A_l)$ we define ωA as follows:

$$\begin{array}{rcl} (\Delta^2)^* & \to & W \in \mathfrak{so}(d) \\ \omega & \mapsto & \omega A \doteq \omega_1 A_1 + \cdots \omega_l A_l; \end{array} \tag{3.4}$$

we define in an analogous way ωQ and ωq.

In the sequel we will need to expand a control $u \in H$ into its Fourier series: we will write

$$u = \sum_{k \in \mathbb{N}_0} U_k \frac{1}{\sqrt{\pi}} \cos kt + V_k \frac{1}{\sqrt{\pi}} \sin kt$$

where $U_k, V_k \in \Delta$; the constant term is zero because of part (a) of Proposition 18 (mean zero condition).

Proposition 19. *Let T_k be the subspace of H with "wave number" k, namely*

$$T_k = \Delta \otimes \operatorname{span}\{\cos kt, \sin kt\}.$$

Then we have the following:

(a) $H = \oplus_{k \geq 1} T_k$, and the sum is orthogonal with respect to the scalar product;

(b) For every $\omega \in (\Delta^2)^*$ we have $\omega Q T_k \subset T_k$ (i. e. each subspace T_k is invariant by ωQ);

(c) Consider the orthonormal basis $\{e_i \otimes \frac{1}{\sqrt{\pi}} \cos kt, e_i \otimes \frac{1}{\sqrt{\pi}} \sin kt\}_{i=1}^d$ for T_k; in this basis the matrix associated to $\omega Q|_{T_k}$ is:

$$\frac{1}{k}(\omega P) \doteq \frac{1}{k} \begin{pmatrix} 0 & \frac{1}{2}\omega A \\ -\frac{1}{2}\omega A & 0 \end{pmatrix}.$$

Proof. Point (a) is just Fourier decomposition theorem; $k \geq 1$ expresses the mean zero condition.
For the other two points, let us consider $u \in T_n$ and $v \in H$, with Fourier series respectively $u = U \frac{1}{\sqrt{\pi}} \cos kt + V \frac{1}{\sqrt{\pi}} \sin kt$ and $v = \sum_{n \geq 1} U_n \frac{1}{\sqrt{\pi}} \cos nt + V_n \frac{1}{\sqrt{\pi}} \sin nt$. By a direct computation we have

$$\begin{aligned}
\langle u, \omega Q v \rangle_H &= \int_I \left\langle \int_0^t U \frac{1}{\sqrt{\pi}} \cos k\tau + V \frac{1}{\sqrt{\pi}} \sin k\tau \, d\tau, \frac{1}{2}\omega A v \right\rangle dt = \\
&= \sum_{n \geq 1} \int_I \frac{1}{k} \left\langle U \frac{1}{\sqrt{\pi}} \sin kt - V \frac{1}{\sqrt{\pi}} \cos kt, \right. \\
&\left. , \frac{1}{2}\omega A \left(U_n \frac{1}{\sqrt{\pi}} \cos nt + V_n \frac{1}{\sqrt{\pi}} \sin nt \right) \right\rangle dt = \\
&= \frac{1}{k} \int_I -\frac{1}{\pi} (\cos kt)^2 \left\langle V, \frac{1}{2}\omega A U_k \right\rangle + \frac{1}{\pi} (\sin kt)^2 \left\langle U, \frac{1}{2}\omega A V_k \right\rangle dt = \\
&= \frac{1}{k} \Big(- \left\langle V, \frac{1}{2}\omega A U_k \right\rangle + \left\langle U, \frac{1}{2}\omega A V_k \right\rangle \Big),
\end{aligned}$$

where the equality between second and third row holds because the only non-zero integrals of products of sines and cosines are

$$\int_I \frac{1}{\pi}(\cos kt)^2\, dt = \int_I \frac{1}{\pi}(\sin kt)^2\, dt = 1.$$

$\square$

Remark 4. We notice that for every $\omega \in (\Delta^2)^*$ the operator ωQ is compact. Indeed, it is the limit of a converging series of operators with finite-dimensional image:

$$S_n = \sum_{i=1}^{n} \omega Q|_{T_i}.$$

Let us prove that the *operator* norm of $\omega Q - \omega S_n$ goes to zero. Given a norm one $v = \sum_{k\geq 1} v_k$, we have:

$$\begin{aligned}\|(\omega Q - \omega S_n)v\|^2 &= \sum_{k\geq n+1} \|\omega Q|_{T_k} v\|^2 \leq \sum_{k\geq n+1} \frac{4\|\omega P\|_{\mathrm{op}}^2 \|v_k\|^2}{k^2} \\ &\leq \frac{4\|\omega P\|_{\mathrm{op}}^2}{(n+1)^2} \sum_{k\geq n+1} \|v_k\|^2 \leq \frac{4\|\omega P\|_{\mathrm{op}}^2}{(n+1)^2}.\end{aligned}$$

In particular, taking square roots, $\|(\omega Q - \omega S_n)v\| \leq \frac{4\|\omega P\|_{\mathrm{op}}}{(n+1)}$, i. e.

$$\|\omega Q - \omega S_n\|_{\mathrm{op}} \to 0$$

We conclude this chapter by describing the spectrum of the operator ωQ. Given ωQ we consider as above the skew-symmetric matrix ωA, which can be put in canonical form as a block matrix of the form

$$\omega A = \mathrm{Diag}\left(\alpha_1(\omega) J_2, \ldots, \alpha_m(\omega) J_2, 0_n\right),$$

where $J_2 \in \mathfrak{so}(2)$ is the standard symplectic matrix and 0_n is the $n \times n$ zero matrix. More precisely on Δ we can find an orthonormal basis $\{X_i, Y_i, Z_j,\ i = 1, \ldots, m,\ j = 1, \ldots, n\}$ for suitable $m, n \in \mathbb{N}$ satisfying $2m+n = d$ such that:

$$\omega A X_i = -\alpha_i(\omega) Y_i, \quad \omega A Y_i = \alpha_i(\omega) X_i, \quad \omega A Z_j = 0.$$

Let us consider now the operator ωQ restricted to T_k. Using the basis for Δ defined above we get the orthogonal basis

$$\left\{ \begin{pmatrix} X_i \\ Y_i \end{pmatrix}, \begin{pmatrix} -Y_i \\ X_i \end{pmatrix}, \begin{pmatrix} X_i \\ -Y_i \end{pmatrix}, \begin{pmatrix} Y_i \\ X_i \end{pmatrix}, \begin{pmatrix} Z_j \\ 0 \end{pmatrix}, \begin{pmatrix} 0 \\ Z_j \end{pmatrix} \right\},$$

for $i = 1, \ldots, m,\ j = 1, \ldots, n$, with eigenvalues $\frac{\alpha_i(\omega)}{k}, \frac{\alpha_i(\omega)}{k}, -\frac{\alpha_i(\omega)}{k}, -\frac{\alpha_i(\omega)}{k}, 0, 0$ respectively. Thus we have proved:

Proposition 20. *The non-zero eigenvalues of the operator ωQ are $\pm\frac{\alpha_i(\omega)}{k}$ with multiplicity two, where $\alpha_i(\omega)$ are the coefficients of the canonical form of ωA and $k \in \mathbb{N}_0$.*

3.1.3 The general structure of geodesics

In this section we fix the final point for the horizontal paths and we define the space of the horizontal curves starting at the origin and ending at p:

$$\Omega_p = F^{-1}(p).$$

Since in the space Ω of horizontal paths we are allowed to compute velocities and their lengths, we define the *Energy* functional:

$$\begin{array}{rccl} J: & \Omega & \longrightarrow & \mathbb{R} \\ & \gamma & \mapsto & \frac{1}{2}\int_0^{2\pi} \|\dot{\gamma}(t)\|^2 dt. \end{array}$$

In the case p is a regular value of the End-point map F, then a critical point of $J|_{F^{-1}(p)}$ is called a *normal geodesic*. From now on the word "geodesic" will simply mean "normal geodesic". We need the regularity of p in order to apply the Lagrange multiplier rule, since we need $F^{-1}(p)$ to be a smooth Hilbert manifold. From now on we will assume p to be a regular value for F; we will see in Proposition 24 for which points this condition is satisfied; it turns out that regular values form a open dense subset of G.

Remark 5. The End-point map F depends on the distribution Δ, and the latter depends on the choice of the skew-symmetric matrices $(A_1, \ldots, A_l)$ that generate the space W. Different choices of basis for W give different distributions: the space of curves Ω_p thus depends on the same choice as a set. However, let us consider the map ϕ introduced in the proof of proposition 16. We denote by $\overline{m}_k^h$ the entries of the inverse matrix M^{-1}, and we call (x, z) the coordinates of the Carnot group associated to the choice of matrices $(A'_1, \ldots, A'_l)$ and (x, y) the coordinates of the Carnot group associated to the choice $(A_1, \ldots, A_l)$.

The isomorphism ϕ is realized as the push forward of the map

$$\begin{cases} x^i = x^i \\ z^h = \sum_{k=1}^l \overline{m}_k^h y^k \end{cases};$$

then we have

$$\begin{aligned} E'_i(x, z) &= \frac{\partial}{\partial x^i} - \frac{1}{2}\sum_{j,h} a'^h_{ij} x^j \frac{\partial}{\partial z^h} = \frac{\partial}{\partial x^i} - \frac{1}{2}\sum_{j,h,k} a'^h_{ij} x^j m_h^k \frac{\partial}{\partial y^k} \\ &= \frac{\partial}{\partial x^i} - \frac{1}{2}\sum_{j,k} a^k_{ij} x^j \frac{\partial}{\partial y^k}. \end{aligned}$$

It follows a natural length-preserving diffeomorphism between the different realizations of the space Ω_p, since through the map ϕ they are diffeomorphic to the same space of control in $L^2(I, \mathbb{R}^d)$.

Using the privileged coordinates described in the subsection 2.1.1, we have that $G \simeq \mathbb{R}^d \oplus \mathbb{R}^l$ we can decompose a covector $\lambda \in T^*G$ as $\lambda = \eta + \omega$, where $\eta \in \Delta^*$ is the "horizontal" part and $\omega \in (\Delta^2)^*$ is the "vertical" one. We have the following proposition.

Proposition 21. *Let u be the control associated to a geodesic with Lagrange multiplier $\lambda \in T^*G$ (i. e. $\lambda d_u F = d_u J$). Then:*

$$u(t) = e^{-(\omega A)t} u_0 \quad \textit{and} \quad 2\eta = (e^{-2\pi\omega A} + \mathbb{1})u_0$$

for a suitable initial vector $u_0 \in \Delta$.

Proof. A simple computation using (3.2) in $\lambda d_u F = u$ gives:

$$u(t) = \eta - (\omega A) \int_0^t u(s) ds + \frac{1}{2} \int_0^{2\pi} (\omega A) u(s) ds.$$

Differentiating the above equation provides $\dot{u} = -(\omega A)u$, which in turn implies $u(t) = e^{-(\omega A)t} u_0$. Substituting the explicit expression $u(t) = e^{-(\omega A)t} u_0$ into the same equation and evaluating at zero gives $u_0 = \eta - \frac{1}{2} u(2\pi) + \frac{1}{2} u_0$. $\square$

If p is not a vertical point and the rank is sufficiently big ($d > l$), then the number of geodesics joining the origin to p is bounded. In order to prove this statement we need a preliminary lemma.

Lemma 22. *For the generic point $p \in G$, the set of normal geodesics starting at the origin and ending at p is discrete.*

Proof. We can parametrize normal geodesics with the vertical part of their initial covector ω and the initial velocity u_0; in this way we obtain a smooth map:

$$f : (\Delta^2)^* \times \Delta \to G$$

defined by:

$$(\omega, u_0) \mapsto \left(\int_I e^{-t\omega A} u_0 dt, \int_I \left\langle \int_0^t e^{-s\omega A} u_0 \, ds, A e^{-t\omega A} u_0 \right\rangle dt \right).$$

If p is a regular value of f (and the set of such p is a residual set) then $f^{-1}(p)$ is a submanifold of $\mathbb{R}^{d+l}$ of dimension zero (possibly non-compact). $\square$

The following Lemma describes the condition for a non-vertical point $p \in G$ to be a critical value for F.

Lemma 23. *The point $p = (p_\Delta, p_\Delta^2) \in G$ is a critical value for F if and only if there is a covector $\lambda \in (\Delta^2)^*$ such that $p_\Delta \in \ker \lambda A$.*

Proof. We compute the differential of F at u:

$$d_u F(v) = \left(\int_I v(t)dt, \frac{1}{2} \int_I \left\langle \int_0^t v(\tau)d\tau, Au(t) \right\rangle + \left\langle \int_0^t u(\tau)d\tau, Av(t) \right\rangle dt \right).$$

We see that the horizontal part is always surjective; then we study the vertical part of the differential $d_u^\perp(v)$ integrating by parts:

$$\begin{aligned} d_u^\perp(v) &= \frac{1}{2} \int_I \left\langle \int_0^t v(\tau)d\tau, Au(t) \right\rangle + \left\langle \int_0^t u(\tau)d\tau, Av(t) \right\rangle dt = \\ &= 2 \int_I \left\langle \int_0^t u(\tau)d\tau, Av(t) \right\rangle dt + \left\langle \int_I u(t)dt, A \int_I v(t)dt \right\rangle = \\ &= -2 \int_I \left\langle A \int_0^t u(\tau)d\tau, v(t) \right\rangle - \langle Ap_\Delta, v(t) \rangle \, dt. \end{aligned}$$

Thus if u is critical for F, there is a covector $\lambda \in (\Delta^2)^*$ such that $\lambda \, d_u^\perp \equiv 0$, namely

$$\lambda A p_\Delta + \lambda A \int_0^t u(\tau)d\tau \equiv 0;$$

setting $t = 0$ we conclude the proof. □

Now we are ready to prove the following:

Proposition 24. *Let G be a Carnot group with a generic Carnot group structure W such that $d > l$. Then for the generic $p \notin \Delta^2$ there is a finite number of geodesics between e and p.*

Proof. We already know by the previous lemma that the set of geodesics (i. e. pairs (ω, u_0) such that $f(\omega, u_0) = p$) is discrete for the generic p. We will exclude that the set of possible Lagrange multipliers is unbounded: this will imply (see below) that the set of initial velocities is bounded as well, hence the set of geodesics ending at p is a discrete set in a compact set, i. e. it is finite.

If ω is a Lagrange multiplier, we can always choose a basis for Δ such that the matrix ωA appears in canonical form: so we get k subspaces of Δ of dimension two on which the matrix ωA is of the form $\alpha_i J_2$ (where J_2 is the standard 2×2 symplectic matrix on the i-th eigenspace and $\alpha_i > 0$), for $i = 1, \dots, k$; for every subspace we take the component u_0^i of the initial velocity u_0. Since the eigenspaces are orthogonal, the computations for the end-point of the geodesic can be performed separately: thus we split the horizontal part of the end-point p_Δ into the components of the eigenspaces of ωA, say p_Δ^i. These components are:

$$p_\Delta^i = \int_0^{2\pi} e^{-t\omega A} u_0^i dt = \left[\frac{1}{\alpha_j} J_2 e^{-t\alpha_i J_2} u_0^i \right]_0^{2\pi} = \frac{1}{\alpha_j} J_2 \left(e^{-2\pi\alpha_j J_2} u_{0j} - u_{0j} \right),$$

on the eigenspaces of ωA. The norm squared of the component p_Δ^i is given by:

$$\begin{aligned}\|p_\Delta^i\|^2 &= \frac{1}{\alpha_i^2} \left\langle J_2 \left(e^{-2\pi\alpha_j J_2} u_0^i - u_0^i \right), J_2 \left(e^{-2\pi\alpha_j J_2} u_0^i - u_0^i \right) \right\rangle = \\ &= \frac{2}{\alpha_i^2} \left(\|u_0^i\|^2 - \left\langle u_0^i, e^{-2\pi\alpha_i J_2} u_0^i \right\rangle \right) = \frac{2}{\alpha_i^2} \|u_0^i\|^2 \left(1 - \cos(2\pi\alpha_i)\right)\end{aligned}$$

The vertical part is more complicated, but we only need to compute its value in the direction of the Lagrange multiplier ω:

$$\begin{aligned}\omega(p_{\Delta^2}) &= \frac{1}{2} \int_0^{2\pi} \left\langle \int_0^t e^{-\omega A s} u_0 ds, \omega A e^{-\omega A t} u_0 \right\rangle dt = \\ &= \frac{1}{2} \int_0^{2\pi} \left\langle \int_0^t -\omega A e^{-\omega A s} u_0 ds, e^{-\omega A t} u_0 \right\rangle dt = \\ &= \frac{1}{2} \int_0^{2\pi} \left\langle e^{-\omega A t} u_0 - u_0, e^{-\omega A t} u_0 \right\rangle dt = \pi \|u_0\|^2 - \frac{1}{2} \left\langle u_0, p_\Delta \right\rangle.\end{aligned}$$

The equation

$$\omega(p_{\Delta^2}) = \pi \|u_0\|^2 - \frac{1}{2} \left\langle u_0, p_\Delta \right\rangle \tag{3.5}$$

is important at first because it tells that it is enough to prove that the set of Lagrange multipliers is bounded. Indeed, we have

$$\|\omega\| \|p_{\Delta^2}\| \geq \pi \|u_0\|^2 - \frac{1}{2} \left\langle u_0, p_\Delta \right\rangle;$$

if the set of Lagrange multipliers is bounded, the set of initial velocities cannot be unbounded otherwise the second term of the inequality would diverge while being limited by a constant.

In order to prove that the set of Lagrange multipliers of the generic end-point is bounded, we suppose on the contrary that we have a sequence ω_n of Lagrange multipliers; this sequence is forced to diverge since the set of the initial data (ω, u_0) for the generic end-point p is discrete. Up to subsequences, we may assume that the normalized Lagrange multipliers $\hat{\omega}_n = \omega_n/\|\omega_n\|$ converge to a covector λ; moreover we can assume the rescaled eigenvalues $\hat{\alpha}_{i,n} = \alpha_{i,n}/\|\omega_n\|$ and the corresponding eigenspaces converge. The first one is true since every eigenvalue is bounded by the norm of the matrix which is 1 by definition; the second one follows from the fact that the set of changes of basis is the orthogonal group which is compact. Hence there exists $c > 0$ such that for every sequence of eigenvalues $\hat{\alpha}_{i,n}$ that doesn't converge $c\|\omega_n\| \le \alpha_{i,n}$.

If we split the second term of the equation (3.5) into its components given by the eigenspaces of ωA we obtain:

$$\omega(p_{\Delta^2}) = \sum_{i=1}^{k} \pi\|u_0^i\|^2 - \frac{1}{2}\left\langle u_0^i, p_\Delta^i \right\rangle.$$

Notice that each term $\pi\|u_0^i\|^2 - \frac{1}{2}\left\langle u_0^i, p_\Delta^i\right\rangle$ is non-negative: in fact $\|p_\Delta^i\| = \left\|\int_0^{2\pi} e^{-\omega A s} u_0^i ds\right\| \le \int_0^{2\pi} \left\|e^{-\omega A s} u_0^i\right\| ds = 2\pi\|u_0^i\|$ and:

$$\frac{1}{2}\left\langle u_0^i, p_\Delta^i\right\rangle \le \frac{1}{2}\|u_0^i\|\|p_\Delta^i\| \le \pi\|u_0^i\|^2.$$

Therefore we have:

$$\|p_{\Delta^2}\| \ge \hat{\omega}_n(p_{\Delta^2}) \ge \pi\frac{\|u_{0,n}^i\|^2}{\|\omega_n\|} - \frac{\left\langle u_{0,n}^i, p_\Delta\right\rangle}{\|\omega_n\|}$$

and if $\|u_{0,n}^i\|$ diverges, the last term of the previous inequality is asymptotic to its first addendum, which turns out to be bounded. Then for the corresponding component of the horizontal part of the end-point we have that:

$$\|p_{\Delta,n}^i\|^2 = \frac{2}{\alpha_{i,n}^2}\|u_{0,n}^i\|^2(1 - \cos(2\pi\alpha_{i,n})) \le 2\frac{\|u_{0,n}^i\|^2}{c^2\|\omega_n\|^2}.$$

If $\|u_{0,n}^i\|$ is bounded the last term of the inequality converges to 0; if it diverges we have that

$$\|p_{\Delta,n}^i\|^2 \le 2\frac{\|u_{0,n}^i\|^2}{c^2\|\omega_n\|^2} \le \text{const.}\frac{1}{\|\omega_n\|}$$

which again converges to 0. So, in the limit, the components of the horizontal part of the end-point are orthogonal to the eigenspaces of the limit matrix λA with non-zero eigenvalues.

Now, if the matrix λA is not singular the horizontal part has to be zero contradicting the hypothesis; thus in order to end the argument, we will prove that the horizontal part of the generic point does not belong to the kernel of any degenerate matrix in the span W.

Let us consider the semi-algebraic set of the couples of degenerate matrices together with the vectors in their kernels, namely

$$X \doteq \{(\eta, x) \in S^{l-1}((\Delta^2)^*) \times \Delta \mid \eta A x = 0,\ \dim\ker \eta A \geq 2\}$$

stratified in subsets with constant dimension of the kernel. Assuming that the Carnot algebra structure defined by W is transversal to every subset of the stratification described in the section 2.3.3 (condition satisfied by the generic W), the stratum W_k of the matrices of k-dimensional kernel has codimension $k\frac{k-1}{2}$. Therefore the semi-algebraic set X has dimension not greater than $\max_{k\geq 2} l - 1 - k\frac{k-1}{2} + k \leq l$. If we project X on Δ we get the set $\tilde{X}$ of points belonging to the kernel of a degenerate matrix of W: this is still a semi-algebraic set and the projection does not increase the dimension. In particular the set

$$Y = \{(p_\Delta, p_{\Delta^2}) \in G \,|\, p_\Delta \in \ker \eta A \quad \text{for some } \eta \neq 0\}$$

has dimension less than $2l$, which is smaller than $d + l = \dim(G)$ under the assumption $d > l$. From Lemma 23 we know the condition above to be the same for p to be regular for F: this concludes the proof. □

Corollary 25. *For the generic rank d Carnot group structure $W \subset \mathfrak{so}(d)$ on $\mathbb{R}^{d+l}$ such that $d > l$, if p is a regular value of the End-point map F and $p \notin \Delta^2$, the homology of $\Omega_p \cap \{J \leq s\}$ stabilizes as $s \to \infty$.*

Proof. If p is a regular value of F, then all geodesics to p are normal and the result follows from the previous proposition. □

Remark 6. In lemma 24 we need the Carnot group structure $W \subset \mathfrak{so}(d)$ to be transversal to the strata of the stratification presented in the section 2.3.3 in order to preserve the coranks of the strata; from now on this will be the meaning of the generic choice of $W \in \mathfrak{so}(d)$.

3.1.4 Geodesics ending at vertical points

In this section we study in more detail the case $p \in \Delta^2$. To start with, we prove that for the generic choice of $p \in \Delta^2$ the set Ω_p is a Hilbert manifold, so that we can still use the Lagrange multiplier rule in order to find the geodesics.

Theorem 26. *For the generic choice of $W \subset \mathfrak{so}(d)$ and a generic $p \in \Delta^2$ the topological space Ω_p is a Hilbert manifold.*

Proof. We will prove that the set of critical values of $q = F|_{F^{-1}(\Delta^2)}$ is contained in a semialgebraic subset of codimension one, from this the conclusion follows.

We first notice that the condition for $p \in \Delta^2$ to be a critical value of q is that there exist $u \in H$ and $\omega \in (\Delta^2)^*$ such that $q(u) = p$ and $dq_u = 2\omega Q u = 0$. In particular if $u = \sum_{k=1}^{\infty} u_k$, where $u_k \in T_k \simeq \mathbb{R}^{2d}$, then $q(u) = \sum_{k=1}^{\infty} \langle u_k, \frac{1}{k}(\omega P) u_k \rangle$ and the condition that $\omega Q u = 0$, by invariance of the spaces T_k, reads $\omega P u_k = 0$ for all $k \geq 1$.

Let us consider the stratification $\mathfrak{so}(d) = \coprod S_r$, where S_r is the set of matrices with constant rank r. Each S_r is smooth and over it we have the smooth bundle $K_r = \{(A, v) \in \mathfrak{so}(d) \times \mathbb{R}^n \mid Av = 0\}$. Since this stratification is homogeneous, then the generic W is transversal to all strata and we have an induced stratification:

$$W = \coprod_{r=0}^{d} W_r, \quad W_r = S_r \cap W.$$

Consider now the vector bundle $K|_{W_r}$ over W_r (the restriction of K_r); notice that for every $\omega A \in W$ we have $\ker(\omega P) = \{(x, y) \in \mathbb{R}^{2d} \mid \omega A x = \omega A y = 0\}$. In particular a smooth section of $K|_{W_r}$ produces also a smooth section of $\{(\omega, z) \in W_r \times \mathbb{R}^{2d} \mid \omega P z = 0\}$ over W_r.

We notice now the following interesting property: if $p \in \Delta^2$ is a critical point for q with Lagrange multiplier $\omega \in W_r$, then $p \in (T_\omega W_r)^\perp$. In fact for every $k \geq 1$ let us consider a smooth curve $\omega(t)$ in W_r with $\omega(0) = \omega$, $\dot{\omega}(0) = \eta \in T_\omega W_j$ and a smooth $z(t) \in \ker(\omega(t)\omega P)$ with $z(0) = u_k$ such that

$$\omega(t)\omega P z(t) = 0,$$

(the existence of such a smooth $z(t)$ follows from the above discussion). Then deriving the above equation we get $\dot{\omega}(0)\omega P z(0) + \omega(0)\omega P \dot{z}(0) = 0$ and considering the scalar product with z gives:

$$\langle z, \dot{\omega}(0)\omega P z \rangle = \dot{\omega}(0)(q(u_k)) = 0$$

which tells $\eta(p) = \sum_{k=1}^{\infty} \eta(q(u_k))$ vanishes for every η in $T_\omega W_r$. We consider now the semialgebraic set:

$$\Sigma_1' = \bigcup_{r=0}^{d} \Big(\bigcup_{\omega \in W_r} (T_\omega W_r)^\perp \Big).$$

Because of the above argument all critical points of q are contained in Σ_1' and we want to show this set is of dimension strictly less than l.

We first check that Σ_1' is indeed semialgebraic (and as a result we compute its dimension); being a finite union, it is enough to prove that each $\bigcup_{\omega \in W_r} (T_\omega W_r)^\perp$ is semialgebraic. To this end consider $T_r = \{(\omega, \eta) \in W_r \times W^* \mid (\eta \in T_\omega W_r)^\perp\}$, which is clearly semialgebraic, and the semialgebraic projection $\pi_2 : T_r \to W^*$ to the second factor. The image of π_2 is semialgebraic and coincides with $\bigcup_{\omega \in W_r} (T_\omega W_r)^\perp$.

Now, each $\bigcup_{\omega \in W_r} (T_\omega W_r)^\perp$ has dimension less than

$$\dim(W_r) + l - \dim(W_r) - 1 \leq l - 1,$$

where the -1 comes from the fact that by homogeneity $(T_{\frac{\omega}{|\omega|}} W_r)^\perp = (T_\omega W_r)^\perp$ (notice in particular that the stratum of maximal dimension is open and produces only the zero, since the orthogonal complement of its tangent space is the zero only).

In particular the dimension of Σ_1' is strictly less than l and the generic $p \in \Delta^2$ is a regular value of q. $\square$

Remark 7. It will be useful for us to set $\Sigma_1 = \overline{\Sigma_1'}$. Since the euclidean closure is smaller than the Zariski closure and dimension of a semialgebraic set is preserved after taking its Zariski closure (which is still a semialgebraic set), then Σ_1 is a *closed* semialgebraic set of codimension one. Points in the complement of Σ_1 are an open dense set of regular values of $q = F|_H$.

Following up the discussion after Proposition 21, we see that in the case the final point of γ_u is in Δ^2, which we know it is equivalent to $\int_I u = 0$, we can apply the Lagrange multiplier rule to the map q. More precisely u is the control associated to a curve which is a geodesic with endpoint $p \in \Delta^2$ if:

$$q(u) = p \quad \text{and there exists } \omega \text{ such that } \omega Q u = u.$$

The covector $\omega \in (\Delta^2)^*$ is the *Lagrange multiplier* associated to u. The complete Lagrange multiplier (i. e. the one arising by using the map F as in Proposition 21), instead of its restriction q), is $\lambda = (u_0, \omega)$. By definition u is a geodesics with Lagrange multiplier ω if and only if:

$$\langle u, v \rangle_H = \langle \omega Q u, v \rangle_H \quad \text{for all } v \in H.$$

Here the final point is not specified, i. e. we are considering all possible geodesics with Lagrange multiplier ω; the final point is recovered by simply applying the expression given in (3.2) to u. We may rewrite the condition above as follows: for all v in $H = \{\int v = 0\}$,

$$\int_0^{2\pi} \langle u(t), v(t) \rangle dt = -\int_0^{2\pi} \langle \omega A U(t), v(t) \rangle dt = 0,$$

where $U(t) = \int_0^t u(s)ds$. The previous condition tells that $u + \omega AU$ is a constant function, or equivalently that $\dot{u} = -\omega Au$. This implies that u must be of the form:

$$u(t) = e^{-t(\omega A)}u_0;$$

moreover $u \in H$, then the condition $\int u = 0$ becomes more explicitly

$$0 = \int_I e^{-t\omega A}u_0 dt = (\omega A)^{-1}\left(u_0 - e^{-2\pi\omega A}u_0\right).$$

It follows that u_0 must be in the *integer* eigenspace of $i\omega A$, i. e.:

$$u_0 = e^{-2\pi\omega A}u_0.$$

We can summarize the above discussion in the following Lemma:

Lemma 27. *Let u be the control associated to a geodesic whose final point is in Δ^2 with Lagrange multiplier ω. Then:*

$$u(t) = e^{-t\omega A}u_0 \quad \textit{with} \quad u_0 = e^{-2\pi\omega A}u_0.$$

Motivated by the previous lemma, for every $\omega \in W$ we define:

$$E(\omega) = \{v \in \Delta \,|\, e^{-2\pi\omega A}v = v,\, v \notin \ker(\omega A)\}.$$

Thus $E(\omega)$ is the set of possible initial data for *non constant* geodesics with Lagrange multiplier ω. In particular we see that in order to have a nonzero initial datum the matrix $i\omega A$ must have nonzero integer eigenvalues, thus the set of all possible Lagrange multipliers coincides with the set:

$$\Lambda = \{\omega \in (\Delta^2)^* \,|\, \det(\omega A - in\mathbb{1}) = 0 \quad \text{for some } n \in \mathbb{N}_0\}.$$

Notice that Λ *is not* an algebraic (or a semialgebraic set): it is indeed given by the infinite union of algebraic sets $\Lambda_n = \{\det(\omega A - in\mathbb{1}) = 0\}$. However Λ is locally algebraic: if we intersect it with a ball, then only a finite number of Λ_n show up.

We discuss now in more detail the structure of the set:

$$E = \{(\omega, v) \in W \times \Delta \,|\, v \in E(\omega)\}.$$

As for Λ, this set *is not* semialgebraic, although if we take the "restriction" $E|_B$ to a compact semialgebraic set B, i. e. we only allow ω to vary on a compact semialgebraic set $B \subset W$, then $E|_B$ becomes semialgebraic.

First for every ω let us consider the canonical skew-symmetric form of ωA:

$$M(\omega)^T(\omega A)M(\omega) = \text{Diag}(\alpha_1(\omega)J_2, \ldots, \alpha_s(\omega)J_2, 0, \ldots, 0)$$

where $M(\omega)$ is an orthogonal matrix, $\alpha_1(\omega), \ldots, \alpha_s(\omega)$ are the *positive non-zero* eigenvalues of $i\omega A$ and $J_2 \in \mathfrak{so}(2)$ is the canonical symplectic matrix. Thus Δ decomposes as the orthogonal sum $\Delta = V_1 \oplus \cdots \oplus V_s \oplus K$, where the V_is are the coordinate two planes and K is the vector space of the last $d-2s$ coordinates. Using this notation we set $V_i(\omega) = M(\omega)V_i$: it is the invariant subspace of ωA associated to the eigenvalue $\alpha_i(\omega)$. In particular we see that:

$$E(\omega) = \bigoplus_{\alpha_i(\omega)\in\mathbb{N}_0} V_i(\omega).$$

Associating to each $v \in E(\omega)$ the control $e^{t\omega A}v$ defines a linear injection of $E(\omega)$ into $H = \oplus_{k\geq 1} T_k$; in particular the previous curve admits the Fourier series decomposition

$$e^{-t\omega A}v = \sum_{k\geq 1} X_k(v)\frac{1}{\sqrt{\pi}}\cos kt - Y_k(v)\frac{1}{\sqrt{\pi}}\sin kt \tag{3.6}$$

and we denote by ϕ_k the linear map $v \mapsto (X_k(v), Y_k(v))$ (the k-th component of the Fourier series of $e^{t\omega A}v$ written in coordinates $T_k \simeq \mathbb{R}^{2d}$).

Let now $v \in V_i(\omega)$ with $\alpha_i(\omega) = k \in \mathbb{N}$; in order to get the expression for $\phi_k(v)$ we compute the Taylor series of $e^{\omega At}v$. We have:

$$\begin{aligned}
e^{-t\omega A}v &= \left(\sum_{n\geq 0}\frac{(\omega A)^n t^n}{n!}\right)v = \\
&= \left(\sum_{m\geq 0}\frac{(\omega A)^{2l}t^{2l}}{(2l)!}\right)v - \left(\sum_{m\geq 0}\frac{(\omega A)^{2l+1}t^{2l+1}}{(2l+1)!}\right)v = \\
&= \left(\sum_{m\geq 0}\frac{(-1)^l k^{2l}t^{2l}}{(2l)!}\right)v - \left(\sum_{m\geq 0}\frac{(-1)^l k^{2l+1}t^{2l+1}}{(2l+1)!}\right)\frac{\omega A}{k}v = \\
&= v\cos kt - \frac{\omega A}{k}v\sin kt
\end{aligned}$$

where in the second line we have used the fact that $(\omega A)^2 v = -k^2 v$ (being $V_i(\omega)$ the space associated to the eigenvalue $\alpha_i(\omega) = k$). This computation implies that:

$$\phi_k(v) = \sqrt{\pi}\begin{pmatrix} v \\ -\frac{\omega A}{k}v \end{pmatrix}.$$

Notice that the same construction can be performed using the linear immersion $v \mapsto e^{t\omega A}v$, which gives the above control with *backward* time; the

Lagrange multiplier for the corresponding geodesic is $-\omega$ and the final point is $-q(e^{-t\omega A}v)$.

Slightly abusing of notation, we will still denote by q the map obtained by composing the endpoint map with the linear immersion $E(\omega) \hookrightarrow H$.

We recall (see Proposition 15) that the Lie algebra $\mathfrak{so}(d)$ of skew symmetric matrices of size d is stratified by the sets $\Gamma_{k|m_1,\ldots,m_r}$ consisting of those matrices A satisfying: $\dim\ker(A) = k$; the numbers $m_1, \ldots, m_r$ are natural non-increasing (they are the multiplicities in the positive spectrum of iA). Each one of these strata is smooth and has codimension $\sum_{i=1}^{r}(m_i^2-1)+\frac{k(k-1)}{2}$. Since (by construction) this stratification is homogeneous, then a generic choice of the Carnot structure $W \subset \mathfrak{so}(d)$ will be transversal to all of the strata and will inherit the stratification (in particular respecting codimensions and smoothness).

To deal with integer eigenvalues we need to refine this stratification, unfortunately ending up with an infinite number of strata, but still with nice properties. More specifically for every $r \geq 0$ we consider $\vec{n} = (n_1, \ldots, n_r) \in \mathbb{N}^r$ with distinct nonzero components and define the *semialgebraic* set

$$\Gamma_{k|m_1,\ldots,m_r|\vec{n}}$$

as follows: we look at the nonzero components of $\vec{n}$, say $n_{j_1}, \ldots, n_{j_\nu}$, and we take those matrices in $\Gamma_{k|m_1,\ldots,m_r}$ such that the eigenvalue with multiplicity m_{j_1} equals n_{j_1}, the one with multiplicity m_{j_2} equals n_{j_2}, and so on.
For example $(0,2,2,0) \in \mathbb{N}^4$ is not an admissible $\vec{n}$ (since there are two equal nonzero entries); on the other hand if $\vec{n} = (0,1,2,0)$ then $\Gamma_{k|m_1,m_2,m_3,m_4|\vec{n}}$ equals the set of all matrices in $\mathfrak{so}(d)$ with multiplicities of the spectrum $\{m_1, m_2, m_3, m_4\}$ *and* with one eigenvalue equal to i (the imaginary unit) with multiplicity m_2 and another equal to $2i$ and with multiplicity m_3.
This operation of fixing some eigenvalues to some integer numbers increases the codimension by ν (the number of nonzero components of $\vec{n}$). Given this new stratification, we will consider from now on the Carnot group structures W which are transversal to all the new strata: even though we restrict the set of W we are considering, the choice remains generic. Then all the strata preserve the codimension, and using the identification $W \cong (\Delta^2)^*$ of Remark 3 we can induce the above stratifications on the set of covectors $(\Delta^2)^*$ as follows:

$$\Lambda_{k|m_1,\ldots,m_r} \doteq \Gamma_{k|m_1,\ldots,m_r} \cap W$$

and

$$\Lambda_{k|m_1,\ldots,m_r|\vec{n}} \doteq \Gamma_{k|m_1,\ldots,m_r|\vec{n}} \cap W.$$

Moreover if we consider admissible $\vec{n}$ only, we stratify the set of Lagrange multipliers Λ as:

$$\Lambda \cong \coprod_{r} \coprod_{\{k,m_1,\dots,m_r\}} \coprod_{\{\vec{n}\in\mathbb{N}^r \text{ admissible}\}} \Lambda_{k|m_1,\dots,m_r|\vec{n}}.$$

As we already noticed, this stratification is not finite even though each stratum is semialgebraic. Nevertheless if we intersect Λ with a compact ball $B \subset W$ only a finite number of the above strata appear and we are locally semialgebraic.

The next lemma tells that for the generic choice of $W \subset \mathfrak{so}(d)$ and a generic $p \in \Delta^2$, the Lagrange multipliers have simple integer spectrum.

Lemma 28. *For a generic Carnot group structure $W \subset \mathfrak{so}(d)$, the generic $p \in \Delta^2$ is not the final point of a geodesic with Lagrange multiplier ω such that ωA has multiple eigenvalues in $i\mathbb{Z}$.*

Proof. First we pick the structure W to be transversal to all strata of the first one of the above stratifications, the one using only the multiplicities in the spectrum (and we know such a property is generic). We stratify now the set Λ by intersecting it with the different strata $\Lambda_{k|m_1,\dots,m_r}$; we are interested only in those strata for which there is at least a multiple integer eigenvalues and we refine the stratification to the above infinite one, by indexing with the admissible $\vec{n} \in \mathbb{N}^r$.

Thus we let $\Lambda_{\{m_j\geq 2\}}$ be one stratum $\Lambda_{k|m_1,\dots,m_r|\vec{n}}$ such that $m_j \geq 2$ for at least one index j with $n_j \neq 0$. Each $\Lambda_{\{m_j\geq 2\}}$ obtained in this way has codimension:

$$\operatorname{codim}_W \Lambda_{\{m_j\geq 2\}} = \sum_{i=1}^{r}(m_i^2 - 1) + \frac{k(k-1)}{2} + \nu.$$

We consider now as above the set $E = \{(\omega, v) \in \Lambda \times \Delta \,|\, e^{2\pi\omega A}v = v,\ v \notin \ker(\omega A)\}$. Over each stratum $\Lambda_{\{m_j\geq 2\}}$ the set $E|_{\Lambda_{\{m_j\geq 2\}}}$ is a smooth vector bundle (it is the restriction to $\Lambda_{\{m_j\geq 2\}}$, which is smooth, of a smooth vector bundle); moreover $E|_{\Lambda_{\{m_j\geq 2\}}}$ is semialgebraic as well (here the vector $\vec{n}$ is fixed).

Consider the smooth map:

$$f : E|_{\Lambda_{\{m_j\geq 2\}}} \to \Delta^2$$

defined by $(\omega, v) \mapsto q(v)$, where $q(v)$ is the final point of the geodesic associated to the control $v(t) = e^{t\omega A}v$. We compute the rank of the differential of

f and show that the assumption $m_j \geq 2$ implies this rank is less than $l-1$; in particular the image of f has measure zero. Since the set of final points of geodesics with Lagrange multipliers with multiple eigenvalues in $i\mathbb{Z}$ is the countable union of the images of the different f obtained as $\vec{n}$ varies over $\mathbb{N}^r$, the result follows.

The differential of f restricted to the base $\Lambda_{\{m_j \geq 2\}}$ has rank smaller than the dimension of $\Lambda_{\{m_j \geq 2\}}$, which is $l - \sum_{i=1}^r (m_i^2 - 1) - \frac{k(k-1)}{2} - \nu$. For the rank of f restricted to the fibers we argue as follows. For every $\omega \in \Lambda_{\{m_j \geq 2\}}$ we consider the invariant subspaces of ωA; for each natural nonzero eigenvalue $\lambda_j(\omega)$ of $i\omega A$ we find an invariant space $V_j(\omega)$ (the real part of the $\lambda_j(\omega)$-eigenspace of $i\omega A$) of dimension $2\mu_j$, twice the multiplicity of $\lambda_j(\omega)$; let's call $I \subset \{1, \ldots, r\}$ the index set for such spaces $V_j(\omega)$ (notice that $I = \{j_1, \ldots, j_\nu\}$).

The restriction of f to each such $V_j(\omega)$ maps μ_j unit circles (lying on distinct orthogonal planes) to a point, in particular the dimension of the kernel of the differential of f on each $V_j(\omega)$ is at least μ_j. Since the dimension of $E(\omega)$ is $2\sum_{j \in I} \mu_j$, we see that the rank of the differential of f on the fibers is at most $\sum_{j \in I} \mu_j$.
In particular we can bound the rank of the differential of f as:

$$
\begin{aligned}
\mathrm{rk}(df) &\leq l - \nu - \sum_{j=1}^r (m_j^2 - 1) - \frac{k(k-1)}{2} + \sum_{j \in I} \mu_j \\
&\leq l - \nu - \sum_{j \in I} (m_j^2 - 1 - m_j) - \frac{k(k-1)}{2} \\
&\leq l - \nu - \sum_{j \in I,\, m_j \geq 2} (m_j^2 - 1 - m_j) - \sum_{j \in I,\, m_j = 1} (m_j^2 - 1 - m_j) \\
&\leq l - \nu - \sum_{j \in I,\, m_j \geq 2} 1 - \sum_{j \in I,\, m_j = 1} (-1) \\
&\leq l - \nu - 1 + (\nu - 1) < l - 1.
\end{aligned}
$$

□

We define now the set $\Sigma_2 \subset \Delta^2$ to be the union of the various

$$f(E|_{\Lambda_{k|m_1,\ldots,m_r,\vec{n}}})$$

where $\vec{n} = (n_1, \ldots n_r) \in \mathbb{N}^r$ is admissible and $m_j \geq 2$ for at least one index j with $n_j \neq 0$. The above lemma says that Σ_2 is the countable union of semialgebraic sets of codimension at least 2 (in particular, for example, Σ_2 has measure zero).

We study now what happens for a $p \in \Delta^2 \backslash (\Sigma_1 \cup \Sigma_2)$ (because of the above argument such p is generic). Every Lagrange multiplier ω of such a p is associated to a matrix ωA with simple spectrum, i. e. ω belongs to a stratum $\Lambda_{k|m_1,\dots,m_r|\vec{n}}$ with all multiplicities equal to 1; for simplicity of notation we omit the string of multiplicities and denote such stratum simply by $\Lambda_{\vec{n}}$. In other words $\Lambda_{\vec{n}}$ is one of the above strata where all eigenvalues are distinct and we have fixed ν of them to be equal to $in_{j_1}, \dots, in_{j_\nu}$ (the nonzero entries of $\vec{n}$).

Lemma 29. *For a generic Carnot group structure $W \subset \mathfrak{so}(d)$, let $\omega \in \Lambda_{\vec{n}}$ and $n_1, \dots, n_\nu$ be the nonzero integer eigenvalues of $i\omega A$ in $\mathbb{N}$; for $j = 1, \dots, \nu$ let also $V_j(\omega)$ be the (two dimensional) invariant subspace of ωA associated to n_j. Then $E(\omega)$ splits as the direct orthogonal sum:*

$$E(\omega) = \bigoplus_{j=1}^{\nu} V_j(\omega).$$

Moreover the image of $q|_{V_j(\omega)}$ is a half line $l_j^+(\omega)$ and:

$$\mathrm{im}(q|_{E(\omega)}) = \mathrm{cone}\{l_1^+(\omega), \dots, l_\nu^+(\omega)\}.$$

Proof. Recall that the space $E(\omega)$ is defined to be $\{v \in \Delta \,|\, e^{-2\pi\omega A}v = v,\ v \notin \ker(\omega A)\}$; the map that associates to a vector $v \in E(\omega)$ the curve $e^{-t\omega A}v$ defines an embedding of $E(\omega)$ into H and if $v \in V_j(\omega)$ then the resulting control must be a linear combination of $\sin(n_j t)$ and $\cos(n_j t)$; in particular $V_{n_j}(\omega) \subset T_{n_j}$. Since the T_k are pairwise orthogonal for each operator $Q_1, \dots, Q_l$, then decomposing $v \in E(\omega)$ into its pieces $v = v_1 + \cdots + v_\nu$ with $v_j \in V_j(\omega)$, we get:

$$q(v) = q(v_1) + \cdots + q(v_\nu),$$

which proves the image of $q|_{E(\omega)}$ is the cone spanned by the vectors $q(V_j(\omega))$ for $j = 1, \dots, \nu$; the orthogonality of the V_j follows from the one of the T_{n_j}.

It remains to prove that the image of $q|_{V_j(\omega)}$ is a half line. By assumption ω belongs to a smooth stratum of codimension ν in W and recalling the definition of $\Lambda_{n_j} = \{\det(\omega A - in_j \mathbb{1}) = 0\}$, we have that:

$$\Lambda_{\vec{n}} = \bigcap_{j=1}^{r} \Lambda_{n_j} \cap \Lambda_{k|1,\dots,1}.$$

The bundle $\coprod_{\omega \in \Lambda_{n_j}} V_j(\omega)$ is smooth (being the restriction of a smooth bundle). In particular for every $\eta \in T_\omega \Lambda_{n_j}$ there are curves $\omega(t) \in \Lambda_{n_j}$ and

$v(t) \in V_j(\omega(t))$ such that $\omega(0) = \omega$, $\dot{\omega}(0) = \eta$ and $v(0) = v$. Deriving the equation $\omega(t)\omega P v(t) = v(t)$ and taking inner product with v we get:

$$0 = \left\langle v, \frac{\dot{\omega}(0)\omega P}{k} v \right\rangle = \eta(q(v))$$

which tells the final point of the geodesic u associated to $e^{\omega A t}v$ is orthogonal to $T_\omega \Lambda_{n_j}$, hence is it contained in a line. On the other hand since ω is the Lagrange multiplier for the geodesic u, we get $\omega(q(u)) = J(u) > 0$, which concludes the proof. □

Everything now is ready for the proof of the Theorem that describes the structure of geodesics.

Theorem 30. *For the generic step 2 Carnot group structure $W \subset \mathfrak{so}(d)$ and a generic $p \in \Delta^2$ the set $\Lambda(p)$ of Lagrange multipliers of geodesic whose final point is p is discrete. Moreover every $\eta \in \Lambda(p)$ belongs to some $\Lambda_{\vec{n}}$ and the set of all geodesics whose endpoint is p with Lagrange multiplier η is a compact manifold of dimension $\nu \leq l$ (ν is the number of nonzero entries of $\vec{n}$) diffeomorphic to the torus* $\underbrace{S^1 \times \cdots \times S^1}_{\nu \text{ times}}$.

Proof. By Lemma 28 we know that for the generic choice of $W \subset \mathfrak{so}(d)$ the generic $p \in \Delta^2$ is a final point only of geodesics with Lagrange multipliers in $\Lambda_{\vec{n}}$ for some $\vec{n} \in \mathbb{N}^{\lfloor d/2 \rfloor}$.
Moreover for every η in $\Lambda(p)$ the set of all geodesics with Lagrange multipliers η and final point p is the preimage of p under the map $q : E|_{\Lambda(p)} \to \Delta^2$.

For every admissible $\vec{n} \in \mathbb{N}^{\lfloor d/2 \rfloor}$ we consider the semialgebraic set $F_{\vec{n}}$ defined by:

$$F_{\vec{n}} = \{(\omega, p) \in \Lambda_{\vec{n}} \times \Delta^2 \,|\, p \in \mathrm{im}(q|_{E(\omega)})\}$$

together with the semialgebraic map $g : F_{\vec{n}} \to \Delta^2$ defined by $(\omega, p) \mapsto p$. Since each point (ω, p) in $F_{\vec{n}}$ has ω in $\Lambda_{\vec{n}}$, then by Lemma 29 the dimension of $F_{\vec{n}}$ is at most l . In fact ω varies on a set of dimension $l - \nu$ and the image of $q|_{E(\omega)}$ is a cone of dimension at most ν. Since $F_{\vec{n}}$ is semialgebraic we stratify it as $F_{\vec{n}} = \coprod_{j=1}^s F_{\vec{n},j}$, where each stratum is smooth semialgebraic of dimension at most l (in fact here the index s depends on $\vec{n}$ as well, but we omit this dependence to simplify notations). Notice that if (ω, p) belongs to a stratum of maximal dimension l, then the cone $q(E(\omega))$ must have maximal dimension ν and p must be in its interior.

The restriction $g_{\vec{n},j} = g|_{F_{\vec{n},j}}$ is smooth semialgebraic, thus by the semialgebraic Sard's lemma (see [8]) the set $C_{\vec{n},j}$ of its critical values is a semialgebraic set of dimension at most $l - 1$. If p is not one of these critical values then $g_{\vec{n},j}^{-1}(p)$ consists of isolated points if $\dim(F_{\vec{n},j}) = l$, and is empty otherwise.

We set Σ_3 to be the union of the critical values of $g_{\vec{n},j}$ ($\vec{n}$ varies over $\mathbb{N}^{\lfloor d/2 \rfloor}$ and j is the stratifying index for $F_{\vec{n}}$ as above); such a union, being a countable union of semialgebraic set of dimension at most $l-1$ has measure zero, hence points belonging to its complement are generic.

On the other hand $\Lambda(p)$ equals the union of the projections on Λ of the various $g_{\vec{n},j}^{-1}(p)$. If we intersect Λ with a compact ball B, we hit only a finite number of strata $\Lambda_{\vec{n}}$ and $\Lambda(p) \cap B$ is discrete; thus for a generic p the set $\Lambda(p)$ is discrete set (possibly infinite).

From Lemma 29 we recall that $q(E(\omega))$ is the cone spanned by the half-lines $l_j^+(\omega) = q(E_j(\omega))$: moreover p is the sum of nonzero vectors belonging to these half-lines, $p = p_1 + \ldots + p_\nu$ with $p_j \in l_j^+(\omega)$. Since the spaces $E_j(\omega)$ are orthogonal with respect to the operators $J, Q_1, \ldots, Q_l$, the condition $q(u) = p$ with $u \in E(\omega)$ can be split up as $q(u_j) = p_j$ with $u_j \in E_j(\omega)$. The condition $q(u_j) = p_j$ is equivalent to $\omega q(u_j) = \omega(p_j)$ since for every covector η orthogonal to p_j the condition $\eta q(u_j) = \eta(p_j)$ is automatically satisfied; on the other hand by the Lagrange Multiplier condition we have that $\omega q(u_j) = J(u_j)$, so that the condition is a positive definite one, $J(u_j) = \omega(p_j)$. This implies that every component $u_j \in E_j(\omega)$ of the geodesic u going to p is constrained on a circle $S^1 \in E_j(\omega)$, from which follows that the critical manifold C_ω is a ν-dimensional torus. $\square$

We can summarize with the aid of a picture in corank 2: we can identify the space of covectors $(\Delta^2)^*$ with its dual Δ^2 which can be seen as the set of vertical points. The Lagrange multipliers are found on the level surfaces of matrices with imaginary integer eigenvalues, namely $\{\omega \mid \det(\omega A - ik\mathbb{1}) = 0\}$, via the identification $(\Delta^2)^* \to W$ given by $\omega \mapsto \omega A$; we call their union Λ. Then we fix the vertical point p and we have to look for the Lagrange multipliers in the half space $\{\omega(p) > 0\} \subset (\Delta^2)^*$. A first type of Lagrange multipliers are found to be the smooth points on the set Λ (so they correspond to matrices with only one imaginary integer eigenvalue) such that the level surface is orthogonal to p: to each one of these Lagrange multipliers it corresponds a critical manifold diffeomorphic to S^1 (figure 3.1).

Then we look for the other Lagrange multipliers in the auto-intersections of the set Λ, namely the points that correspond to matrices with more than 1 imaginary integer eigenvalue. For instance let ω correspond to a matrix with imaginary integer eigenvalues $k_1, \ldots, k_\nu$. For every k_j we consider the direction orthogonal to the hypersurface $\{\omega | \det(\omega A - ik_j\mathbb{1}) = 0\}$ and we choose the half-line of vectors with positive scalar product with p, namely by abuse of notation $\langle l_j^+(\omega), p \rangle > 0$. We consider the cone spanned by these directions $\mathrm{cone}\{l_1^+(\omega), \ldots, l_\nu^+(\omega)\}$: ω is a Lagrange multiplier if and only if p belongs to that cone (figure 3.2).

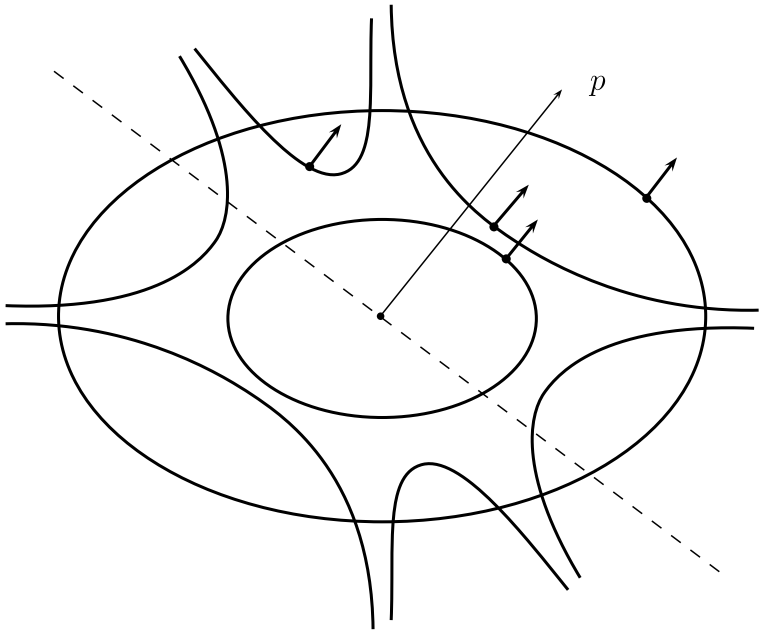

Figure 3.1: Lagrange multipliers of the first type.

The type of critical manifold depends on the number of imaginary integer eigenvalues: indeed it is a copy of $\underbrace{S^1 \times \cdots \times S^1}_{\nu \text{ times}}$. Moreover the Energy of the critical manifold C_ω corresponding to the Lagrange multiplier ω is quite easy to compute: it is indeed $J(C_\omega) = \langle \omega, p \rangle$.

Up to now the assumptions for p come from Theorem 26, Lemma 28 and Theorem 30: specifically we require $p \notin \Sigma_1 \cup \Sigma_2 \cup \Sigma_3$ (where Σ_3 is defined in the proof of Theorem 30).

Remark 8 (Genericity of the vertical point). The set $\Sigma_1 \cup \Sigma_2 \cup \Sigma_3$ is the countable union of semialgebraic closed sets of codimension greater or equal than 1, from which it follows that it has measure zero. Moreover, as we already noticed, every finite radius ball in Δ^2 intersects a finite number of these semialgebraic sets; it follows that the complement is an open dense set.

3.1.5 Non-degenerate critical manifolds

Since critical points of the Energy functional J are not isolated (they arrange themselves into compact manifolds) we cannot apply Morse Theory in its standard version. We need to apply *Morse-Bott Theory*, introduced in section 2.2: we recall that it allows to prove the same results as for the ordinary theory if in the definitions non-degenerate critical points are replaced by

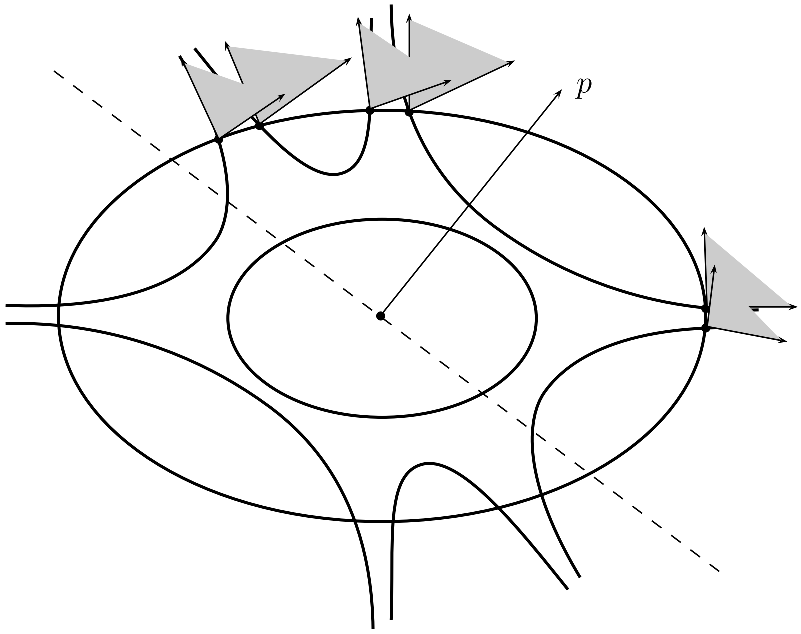

Figure 3.2: Lagrange multipliers of the second type.

nondegenerate critical manifolds .

Theorem 31. *For the generic step 2 Carnot Group structure $W \subset \mathfrak{so}(d)$ and a generic point $p \in \Delta^2$, the Energy functional J restricted to Ω_p is a Morse-Bott function.*

Proof. We assume all the genericity conditions of the previous theorems to be satisfied. We need to prove that we can possibly restrict the set of "good" final points p to a smaller (but still dense) set for which $J|_{\Omega_p}$ is Morse-Bott.

Part (a) of the definition of a Morse-Bott function given in the subsection 2.2, immediately follows from Theorem 30.

We proceed to prove part (b). Let us take a Lagrange multiplier η, its corresponding critical manifold C_η and a point $u \in C_\eta$. Since we already know that the critical manifold is compact, it remains to show that the Hessian of the energy J is non-degenerate outside the tangent space to the critical manifold C_η. Since both q and J are quadratic, they coincide with their second derivatives. By the Lagrange multiplier rule we get the expression for the Hessian:

$$\mathrm{He}_u(q) = (d^2 J - \eta D^2 q_u)|_{T_u \Omega_p} = (J - \eta q)|_{T_u \Omega_p},$$

where we have $T_u \Omega_p = \ker D_u q$. Notice that to the quadratic form defined by the Hessian it corresponds the self-adjoint operator $\mathbb{1} - \eta Q$.

The Hessian is degenerate in the direction of $v \in T_u\Omega_p$ if and only if:

$$\langle v - \eta Q v, x\rangle = 0 \qquad \forall x \in T_u\Omega_p,$$

meaning that $v - \eta Qv$ is orthogonal to $T_u\Omega_p$. Since the tangent space $T_u\Omega_p$ is the orthogonal space to $\mathrm{span}\{Q_1u, \ldots, Q_lu\}$, then $v - \eta Qv$ is a linear combination of the vectors Q_iu, namely $\lambda_1 Q_1 u + \ldots + \lambda_l Q_l u$.

We can eventually restate the degeneracy condition by the following equations:

$$\begin{cases} v - \eta Q v = \lambda Q u \\ \langle v, Qu\rangle = 0 \end{cases} \tag{3.7}$$

where $\lambda = (\lambda_1, \ldots, \lambda_l) \in (\Delta^2)^*$ as above.

If we take $\lambda = 0$ we see that the degeneracy condition is satisfied by the vectors in:

$$E(\eta) \cap T_u\Omega_p = T_u C_\eta,$$

and we have to prove that for the generic choice of p the degeneracy equation (3.7) does not admit other solutions.

Let us consider the smooth manifold $\Lambda_{\vec{n}}$ of Lagrange multipliers containing η (the definition of $\Lambda_{\vec{n}}$ is given before Lemma 29). Let us call as before $E_{\vec{n}}$ the fiber bundle with base space $\Lambda_{\vec{n}}$ and fiber $E(\omega)$ with $\omega \in \Lambda_{\vec{n}}$ (see the above discussion).

The tangent space to $E_{\vec{n}}$ at (u, η) is determined as follows: take a curve $(u(t), \eta(t))$ in $E_{\vec{n}}$ based on (u, η) and compute its tangent vector in $t = 0$. Differentiating the condition $\eta(t)Qu(t) = u(t)$, we get:

$$\dot{u} - \eta Q \dot{u} = \dot{\eta} Q u,$$

which is the same condition as for the degeneracy of the Hessian (the first equation in (3.7)).

We consider now the smooth semialgebraic map:

$$f : E_{\vec{n}} \to \Delta^2 \quad \text{given by} \quad (\omega, v) \mapsto q(v).$$

The set of regular values of f is a dense subset of Δ^2 (it is the complement of a semialgebraic set of codimension at least one): this subset is the good one we want to restrict to. In other words we consider Σ_4 to be the union of the set of critical values of the various $f : E_{\vec{n}} \to \Delta^2$ as $\vec{n}$; the complement of Σ_4 contains generic points[1]. On the preimage of a "good" p we know that the differential of f is surjective with rank l; moreover:

$$\dim E_{\vec{n}} = \dim \Lambda_{\vec{n}} + \dim E(\eta) = l - \nu + 2\nu = l + \nu.$$

[1]Thus at this stage $p \in \Delta^2 \backslash (\Sigma_1 \cup \Sigma_2 \cup \Sigma_3 \cup \Sigma_4)$.

Looking at the dimensions of the domain and the range of $d_u f$ we get that the kernel has dimension ν. On one hand, the kernel of $d_u f$ is the vector space satisfying both the equations (3.7) for the degeneracy of the Hessian; on the other hand we have the inclusion:

$$E(\eta) \cap T_u\Omega_p \subset \ker d_u f.$$

Since both these spaces have dimension ν, they must be equal. It follows that with all the above generic restrictions on p, the only directions of degeneracy for the Hessian are in $T_u C_\eta$.

Remark 9. We know that the operator ηQ is compact, and that the eigenvalues are of the form $\pm\frac{\alpha_i}{k}$ with $k \in \mathbb{N}$ non-zero and $i = 1, \ldots, s < \infty$. Then the eigenvalues of the Hessian of the energy on a critical point are $1 \pm \frac{\alpha_i}{k}$; it follows that the number of negative eigenvalues is always finite, so that the index of every critical manifold is finite.

It remains to prove property (c) of the definition (the Palais-Smale property). Let us consider a sequence $\{u_k\}$ in Ω_p with energy $\|u_k\|^2$ bounded by E and such that $\nabla\psi_{u_k} \to 0$, where $\psi \doteq J|_{\Omega_p}$.

The gradient $\nabla\psi_u$ is the orthogonal projection of $\nabla J_u = u$ on the space $\operatorname{span}\{Q_1 u, \ldots, Q_l u\}^\perp$, then if we define π_u to be the orthogonal projection on the space $\operatorname{span}\{Q_1 u, \ldots, Q_l u\}$ we have

$$\nabla\psi_u = u - \pi_u u\,.$$

From $\{u_k\}$ we can extract a subsequence (we keep calling it $\{u_k\}$) such that $\|u_k\|^2 \to L$. Now we can compute

$$\|\nabla\psi_u\|^2 = \langle u, u\rangle - 2\langle u, \pi_u u\rangle + \langle \pi_u u, \pi_u u\rangle = \langle u, u\rangle - \langle \pi_u u, \pi_u u\rangle\,,$$

where the second equality follows from $\langle u, \pi_u u\rangle = \langle \pi_u u, \pi_u u\rangle$ being π_u an orthogonal projection.

Since $\langle u_k, u_k\rangle \to L$, and $\|\nabla\psi_{u_k}\|^2 = \langle u_k, u_k\rangle - \langle \pi_{u_k} u_k, \pi_{u_k} u_k\rangle \to 0$, it follows that

$$\langle \pi_{u_k} u_k, \pi_{u_k} u_k\rangle \to L\,.$$

Now, every $\pi_{u_k} u_k$ is a linear combination of the vectors $Q_i u_k$, namely

$$\pi_{u_k} u_k = \sum_{i=1}^{l} \eta_k^i Q_i u_k\,,$$

and by computing its norm we get

$$\sum_{i,j=1}^{l} \eta_k^i \eta_k^j \langle Q_i u_k, Q_j u_k\rangle \to L\,. \tag{3.8}$$

Since the operators Q_i are compact, they map the bounded sequence $\{u_k\}$ to a sequence $\{Q_i u_k\}$ with limit points, so we can iteratively extract converging subsequences (again we keep calling them $\{u_k\}$) and we have $Q_i u_k \to v_i$. In this way the equation (3.8) becomes

$$\sum_{i,j=1}^{l} \eta_k^i \eta_k^j \langle v_i, v_j \rangle \to L\,,$$

where the coefficients $\langle v_i, v_j \rangle$ give the scalar product of the whole Hilbert space H restricted to the finite dimensional subspace $V = \mathrm{span}\{v_i, \dots, v_l\}$. Now we have a bounded sequence of vectors $\eta_n = \sum_{i=1}^{l} \eta_n^i v_i$ in $\mathbb{R}^l$ from which we can extract a converging sequence with limit η.

So the sequence $\pi_{u_k} u_k = \sum_{i=1}^{l} \eta_k^i Q_i u_k$ tends to $v = \sum_{i=1}^{l} \eta^i v_i$, and since

$$0 = \lim_{k\to\infty} \|\nabla \psi_{u_k}\|^2 = \lim_{k\to\infty} \|u_k - \pi_{u_k} u_k\|^2 = \lim_{k\to\infty} \|u_k - v\|^2\,,$$

also the sequence $\{u_k\}$ tends to v. Moreover, since $v_i = \lim_{k\to\infty} Q_i u_k = Q_i v$ we have that $v = \eta Q v$ which is the condition for v to be a critical point of the Energy. $\square$

Remark 10 (Genericity of the vertical points, revisited). In the previous proof we have restricted the final vertical point to be outside a new set Σ_4 in order to have the Energy J satisfying the Morse-Bott conditions. From the semialgebraic Sard's Lemma it follows that Σ_4 is a countable union of algebraic subsets of measure 0, so Σ_4 has measure 0 as well. This is the last condition we will require on the points, and from now on the vertical points are assumed to be in $\Delta^2 \backslash (\Sigma_1 \cup \Sigma_2 \cup \Sigma_3 \cup \Sigma_4)$; still, such a choice is generic.

3.2 The topology of horizontal-path spaces

3.2.1 Paths with bounded Energy

Despite Theorem 26 shows that some of the Ω_p may not be Hilbert manifolds, they are in fact all homotopy equivalent to each other; the argument is a simple modification of the standard one for loop spaces and appeared first in [14]; we recall it here for convenience of the reader.

Theorem 32. *For every $p_1, p_2 \in G$ the spaces Ω_{p_1} and Ω_{p_2} are homotopy equivalent.*

Proof. It is sufficient to prove that for every $p \in G$ the space Ω_p is homotopy equivalent to Ω_e. To this end let $\gamma_0 \in \Omega_p$ be a fixed *horizontal* path and define the map:

$$A : \Omega_e \to \Omega_p$$

by concatenation of loops in Ω_e with γ: $A(\gamma) = \gamma_0\gamma$ (velocities have to be rescaled). Let also $\hat{\gamma}_0$ be γ_0 with backward time (it connects p to e); then define

$$B : \Omega_p \to \Omega_e$$

by concatenation with $\hat{\gamma}_0 : B(\gamma) = \hat{\gamma}_0\gamma$. Let now γ_ϵ, $\epsilon \in [0, 1]$ be the paths:

$$\gamma_\epsilon(t) = \gamma_0(\epsilon(1 - t))$$

and $L_\epsilon : \gamma \to \gamma_\epsilon\hat{\gamma}_\epsilon\gamma$. The maps L_ϵ give a homotopy between the identity $L_0 = \mathrm{id} : \Omega_e \to \Omega_e$ and $L_1 = AB$. In a similar way BA is homotopic to the identity on Ω_p and the two spaces are homotopy equivalent. □

As a corollary we see that Ω_p is contractible (in particular all its nonzero Betti numbers vanish).

Corollary 33. *For every $p \in G$ the topological space Ω_p is contractible.*

Proof. By the above Theorem it is enough to show that Ω_e is contractible, and this is obvious since it is given by *homogeneous* equations. □

Even though the space of horizontal paths Ω_p is contractible, there is much more to say about its topology. We studied in the previous section the structure of geodesics, and we saw that whenever p is vertical there is an infinite number of geodesics; every non-degenerate critical manifold contains indeed an infinite number of geodesics, but most importantly the critical manifolds are countable. This may suggest thay the topology of Ω_p is not so trivial, perhaps from a different point of view.

Now we focus on the object

$$\Omega_p^s \doteq \Omega_p \cap \{J \leq s\};$$

they are the sublevels of the Energy, and by letting the Energy grow we study how the topology of the sublevels change, finding very interesting properties. Moreover it will come out when we will study the horizontal paths in the case of general sub-Riemannian manifolds.

Proposition 24 tells us that in the case p is not a vertical point the number of geodesics joining e to p is finite; in particular if $s > 0$ is large enough f does not have critical points on $\{f \geq s\}$ and the topology of Ω_p^s stabilizes.

Proposition 34. *If p is not a vertical point, then for every s large enough and $t > 0$ the inclusion:*

$$\Omega_p^s \hookrightarrow \Omega_p^{s+t} \quad \textit{is a homotopy equivalence.}$$

As we already mentioned we focus on the case $p \in \Delta^2$, and we *will not* use Morse-Bott theory to give a lower bound on $b(\Omega_p^s)$ by counting critical manifolds; we will instead use it to reduce the problem to the study of intersection of real quadrics. In fact the following proposition shows that Ω_p^s is homotopy equivalent to its boundary $\partial\Omega_p^s$; since the map q is quadratic, the space $\partial\Omega_p^s$ is an intersection of infinite-dimensional quadrics (it is given by the quadratic equations $\|u\|^2 = 2s$ and $q(u) = p$).

Proposition 35. *For a generic choice of the Carnot group structure $W \subset \mathfrak{so}(d)$ and a generic point $p \in \Delta^2$, for almost every s the following isomorphism holds:*

$$H_*(\Omega_p^s) \simeq H_*(\partial\Omega_p^s).$$

Proof. We first notice that the generic s is not a critical value for the energy. Let us consider now the Morse-Bott function $g = -J$ and let us denote by X^a the set $\{g \leq a\}$. The critical manifolds of g are the same as for J, except that the *index* of each one of them for g is infinite (since these manifolds have *finite* index for J, then they must have *infinite* index for g).

After passing a critical value c with corresponding critical manifold C, the relative homology of the Lebesgue set is given by (i.e. "the homology changes by"):

$$H_*(X^{c+\delta}, X^{c-\delta}) \simeq H_*(D_C^-, \partial D_C^-).$$

We recall that D_C^- is the unit disk bundle in the fiber bundle over C on which the Hessian of the Morse-Bott function is negative definite; see Theorem 5 and the subsequent discussion from Appendix 2.2. Notice that here the choice of the coefficients field $\mathbb{Z}_2$ prevents us from the problem of orientability of this bundle.

This relative homology is zero: since the index of C is infinite, then both D_c^- and ∂D_c^- retract on C (this follows from the fact that the infinite dimensional sphere is contractible).

We can conclude our proof by observing that even though we pass critical values for $-J$, the homology remains the same of $\partial\Omega_p^s$ until we get the whole Ω_p^s. $\square$

Thus we see that, being Ω_p contractible, each of the Betti numbers $b_i(\Omega_p^s)$ ($i > 0$) eventually vanishes as $s \to \infty$. Despite this their sum can still grow: the smallest $i > 0$ for which $b_i(\Omega_p^s) \neq 0$ will get bigger and bigger and the amount of topology can increase as well: we are interested in understanding quantitatively this phenomenon.

3.2.2 Asymptotic Morse-Bott inequalities

Before giving an explicit bound to $b(\Omega_p^s)$, we will see what *would* this bound *be* if we were to use Morse-Bott inequalities only. The Morse-Bott inequalities bound will follow from the count of the muber of critical manifolds with energy less than s. It turns out that this bound is much worse than the actual one. In fact one has:

$$\text{Card}\{\text{critical manifolds with energy less than } s\} = O(s)^l$$

against the actual bound $b(\Omega_p^s) = O\,(s)^{l-1}$ (this will be proved in the next section).

The following proposition will be fundamental for the sequel: essentially it allows to turn the direct limits arguments into a quantitative form. Roughly it says that the wave numbers of the controls associated to the geodesics grow at most with the order of their energy.

To deal with these ideas, we introduce the following useful notation: for every $L \in \mathbb{N}$ we define the space of controls with wave number at most L:

$$T^L \doteq \bigoplus_{k \leq L} T_k,$$

where the spaces T_k defined in Proposition 19 are

$$T_k = \Delta \otimes \mathit{span}\{\cos kt, \sin kt\}.$$

Proposition 36. *For the generic choice of the Carnot group structure $W \subset \mathfrak{so}(d)$ and the generic point $p \in \Delta^2$ there exists a constant $c_p > 0$ such that for every geodesic $\gamma \in \Omega_p \cap \{J \leq s\}$, its associated control belongs to $T^{\lfloor sc_p \rfloor}$.*

In order to prove the previous proposition we first need the following lemma:

Lemma 37. *For the generic choice of the Carnot group structure $W \subset \mathfrak{so}(d)$ and the generic point $p \in \Delta^2$ there exists a constant $c_p > 0$ such that for every Lagrange multiplier ω associated to p, the following inequality holds:*

$$\frac{\langle \omega, p \rangle}{\|\omega\|} \geq \frac{1}{c_p}.$$

Remark 11. Being the quantity $\langle \omega, p \rangle / \|\omega\|$ the cosine between ω and p times the norm of p, the lemma says that the Lagrange multipliers for p are contained in a convex acute cone in W.

The norm on the space of covectors $(\Delta^2)^* \cong W$ is the one induced by the inclusion $W \hookrightarrow \mathfrak{so}(d)$ where $\langle X, Y \rangle = \mathit{Trace}(X^T Y)$.

Before giving the proof of the lemma we show how it implies Proposition 36.

Proof of proposition 36. Lemma 37 is equivalent to

$$\|\omega\| \leq c_p\langle \omega, p\rangle.$$

The norm of $\omega A \in W$ can be written in terms of its eigenvalues $\alpha_1, \ldots, \alpha_s$: $\|\omega\| = \|\omega A\| = \sqrt{2\alpha_1^2 + \ldots + 2\alpha_s^2}$; it follows that every eigenvalue of ωA is smaller than the norm of ωA. Since ω is a Lagrange multiplier, if $u = u_{k_1} + \ldots + u_{k_l}$ then the k_j are integer eigenvalues of $i\omega A$. Since the energy of a geodesic u associated to ω is $J(u) = \langle \omega, p\rangle$, we have

$$k_j \leq \|\omega\| \leq c_p\langle \omega, p\rangle \leq c_p s. \tag{3.9}$$

□

Now we go back to the proof of Lemma 37.

Proof of Lemma 37. Suppose on the contrary that the constant bounding $\langle \omega, p\rangle/\|\omega\|$ from below doesn't exist, so that we can find a sequence of Lagrange multipliers ω_n such that, setting $\hat{\omega}_n \doteq \omega_n/\|\omega_n\|$, the sequence $\hat{\omega}_n(p) \to 0$. Since the sequence $\hat{\omega}_n$ is contained in S^{l-1} which is compact, we can assume (up to subsequences) that it converges, with limit λ such that $\langle \lambda, p\rangle = 0$ by hypothesis. Up to subsequences we can also assume that every Lagrange multiplier ω_n has the same number of integer eigenvalues (all distinct by Lemma 28), say ν. For every Lagrange multiplier ω_n we have the cone of the endpoints of the geodesics associated to ω_n; the Lagrange multiplier ω_n is contained in the intersection of ν hypersurfaces of matrices with constant eigenvalue equal the imaginary integers $ik_1(n), \ldots, ik_r(n)$. The direct sum $E(\omega)$ of the associated eigenvalues contains all the geodesics with Lagrange multiplier ω and $q(E(\omega))$ is the cone spanned by the normal vectors to each of these ν surfaces (Proposition 29), where the normal vector has to be chosen with positive scalar product with ω_n (since the Energy is positive).

Let us call l_j^+ the normal vectors to the surfaces with eigenvalue equal to $k_j(n)$ respectively; up to subsequences again we can assume that every direction $l_j^+(n)$ converges to some l_j^+. The point p is contained in the interior of every cone (Theorem 30) and it can be written as $p = \sum_{j=1}^{\nu} c_j(n) l_j^+(n)$ with $c_j(n) > 0$ for every j and every n. Since $\langle \lambda, p\rangle = 0$, we have

$$0 = \langle \lambda, p\rangle = \sum_{j=1}^{\nu}\langle \lambda, l_j^+\rangle = \lim_{n\to\infty}\sum_{j=1}^{\nu} c_j(n)\langle \hat{\omega}_n, l_j^+(n)\rangle;$$

the terms of the sum above must converge to 0 one by one because they are non-negative. Not every $c_j(n)$ can converge to 0, otherwise p would be 0 as well. Therefore at least one of the terms $\langle \hat{\omega}_n, l_j^+(n) \rangle$ converges to 0; in the limit p is a linear combination of directions l_j^+ with $j \in I$ for a set of indexes I such that $\langle \lambda, l_j^+ \rangle = 0$ for every $j \in I$.

Now we are going to see what happens to Lagrange multipliers associated to directions l_j^+ such that $\lim_{n\to\infty} \langle \hat{\omega}_n, l_j^+(n) \rangle = 0$. Let us take the smooth hypersurface $S_\alpha \subset W$ of the matrices with an eigenvalue equal to a given $i\alpha \in i\mathbb{R}$ and all the other eigenvalues different from $i\alpha$. Then we take a point $\omega A \in S_\alpha$ and we compute the angle between ωA and the normal vector to the surface S_α at ωA. The surface S_α can be given as a zero locus of the real valued function $s_\alpha(\eta) \doteq \det(\eta A - i\alpha \mathbb{1})$ in the even-dimensional case and $i \det(\eta A - i\alpha \mathbb{1})$ in the odd-dimensional case (here for simplicity we discuss the case d is even, but the proof for the odd-dimensional case is analogous).

We may assume that the matrices $A^1, \ldots, A^l$ form an *orthonormal* basis for W; moreover we can choose an orthonormal basis for Δ such that the matrix ωA is written in canonical form, i.e. $\omega A = \mathrm{Diag}(\alpha J_2, \alpha_2 J_2, \ldots, \alpha_{d/2} J_2)$. Now we compute the differential of s_α at ωA:

$$(ds_\alpha)_{\omega A} = \sum_{i=1,j=1}^{d} \sum_{k=1}^{l} \frac{\partial s_\alpha}{\partial m_{ij}} \frac{\partial m_{ij}}{\partial \eta_k} d\eta_k = \sum_{i=1,j=1}^{d} \sum_{k=1}^{l} \mathrm{adj}(\omega A - i\alpha \mathbb{1})^{ij} a_{ij}^k \, d\eta_k$$

where m_{ij} are the variables for the entries of the matrices, η_k are the coordinates on W given by the components of the covectors in $(\Delta^2)^*$, $\mathrm{adj}(\omega A - i\alpha \mathbb{1})^{ij}$ is the ij entry of the adjugate matrix of $\omega A - i\alpha \mathbb{1}$ and a_{ij}^k are the entries of the matrix A^k. The matrix $\omega A - i\alpha \mathbb{1}$ takes the form

$$\mathrm{Diag}\left(\begin{pmatrix} -i\alpha & \alpha \\ -\alpha & -i\alpha \end{pmatrix}, \begin{pmatrix} -i\alpha & \alpha_2 \\ -\alpha_2 & -i\alpha \end{pmatrix}, \ldots, \begin{pmatrix} -i\alpha & \alpha_{d/2} \\ -\alpha_{d/2} & -i\alpha \end{pmatrix} \right);$$

setting $\beta \doteq \prod_{i=2}^{d/2} (a_i^2 - a^2)$, the adjugate matrix is

$$\mathrm{adj}(\omega A - i\alpha \mathbb{1}) = \mathrm{Diag}\left(\begin{pmatrix} -i\alpha\beta & \alpha\beta \\ -\alpha\beta & -i\alpha\beta \end{pmatrix}, \begin{pmatrix} 0 & 0 \\ 0 & 0 \end{pmatrix}, \ldots, \begin{pmatrix} 0 & 0 \\ 0 & 0 \end{pmatrix} \right),$$

and so we get

$$(ds_\alpha)_{\omega A} = 2\alpha\beta a_{12}^k \, d\eta_k.$$

Therefore we have

$$\langle (ds_\alpha)_{\omega A}, \omega A \rangle = 2\alpha\beta\omega_k a_{12}^k = 2\alpha^2\beta.$$

Since the basis $(A^1, \ldots, A^l)$ is orthonormal for W and so is $(d\lambda_1, \ldots, d\lambda_l)$ for W^*, the norm $\|(ds_\alpha)_{\omega A}\|$ is easily computed:

$$\|(ds_\alpha)_{\omega A}\| = 2|\alpha\beta|\sqrt{(a^1_{12})^2 + \ldots + (a^l_{12})^2} = 2|\alpha\beta|\|a_{12}\|$$

where $a_{12} \doteq (a^1_{12}, \ldots, a^l_{12})$. Now we can compute the cosine of the angle θ between ω and the normal to S_α at ωA:

$$\cos\theta = \frac{\langle (ds_\alpha)_{\omega A}, \omega A\rangle}{\|(ds_\alpha)_{\omega A}\| \, \|\omega A\|} = \frac{2\alpha^2\beta}{2|\alpha\beta| \, \|a_{12}\| \, \|\omega\|} = \pm\frac{\alpha}{\|a_{12}\| \, \|\omega\|}.$$

Let us go back to the directions l^+_j such that

$$\langle \lambda, l^+_j\rangle = \lim_{n\to\infty} \langle \hat{\omega}_n, l^+_j(n)\rangle = 0;$$

recall that the imaginary integers $ik_j(n)$ are the eigenvalue corresponding to the eigenspaces mapped on $\mathbb{R}l^+_j(n)$. Notice that the constant $\|a_{12}\|$ in the previous computation depends on the space W and on the basis we choose in order to put ωA in canonical form, but now the basis is a priori different for every $\omega_n A$; indeed we have that $\|a_{12}\| \leq \|A\|$ where the norm $\|A\|$ is the one induced by the norm on Δ on its exterior algebra. It follows that

$$|\langle \hat{\omega}_n, l^+_j(n)\rangle| \geq \left| \frac{\alpha_i(n)}{\|A\| \, \|\omega_n\|} \right|,$$

and since the first term converges to 0, so does the second one.

Notice that the vectors $l^+_j(c)$, orthogonal to the surfaces with eigenvalue equal to $k_i(n)$, are orthogonal to the surfaces with eigenvalue equal to $\frac{k_i(n)}{\|\omega_n\|}$ as well, since the surfaces are the same up to homothety. The point p is a linear combination of the directions l^+_j with $j \in I$ with the property that $\langle \lambda, l^+_j\rangle = 0$: these directions are the images of the limit eigenspaces $E_j(\lambda)$ of the sequences $E_j(\hat{\omega_n})$. Since their eigenvalues are $\frac{k_j(n)}{\|\omega_n\|}$ which converge to 0, the eigenvalues of the eigenspaces $E_j(\lambda)$ for $j \in I$ are 0. It follows that p belongs to the image by q of the kernel of ωQ, which means that p is a critical value for q, and this is absurd by hypothesis. $\square$

We prove now the theorem that gives an upper bound for the number of critical manifolds with energy bounded by s.

Theorem 38. *For the generic choice of step 2 Carnot group structure $W \subset \mathfrak{so}(d)$ and of $p \in \Delta^2$:*

$$\mathrm{Card}\{\textit{critical manifolds with energy less than } s\} = O(s)^l.$$

Proof. Let us start by considering the vector space $P_{l,d}$ of real polynomials of degree d in l variables. For every $\nu = 1, \ldots, l$ we set $f = (f_1, \ldots, f_\nu) \in P_{l,d}^\nu$ and:

$$\begin{aligned} Y_\nu =&\{(f, V, \omega) \in P_{l,d}^\nu \times G(l-\nu, l) \times \mathbb{R}^l \mid \\ &\omega \in Z(f) \backslash \mathrm{Sing}(Z(f)),\, V = (T_\omega Z(f))^\perp,\, p \in V\}. \end{aligned}$$

Since Y_ν is semialgebraic, we can consider the semialgebraic projection to the first factor $\pi : Y_\nu \to P_{l,d}^\nu$ and stratify $P_{l,d}^\nu = \coprod_{j=1}^s P_j$ such that π is semialgebraically trivial over each stratum (see [8]). In particular since there are finitely many strata, there exists a number β_ν such that for every $f \in P_{l,d}^\nu$:

$$b_0(\pi^{-1}(f)) \leq \beta_\nu$$

(only a finite number of fibers, up to semialgebraic homeomorphism, appear). In particular, in the case when there are finitely many ω with normals $V = (T_\omega Z(f))^\perp$ containing p, this construction implies their number is bounded by β_ν.

Consider now for every $n \in \mathbb{N}$ the polynomial $f_n \in P_{l,d}$ defined by:

$$f_n(\omega) = i^d \det(\omega A - in\mathbb{1})$$

(the i^d factor has the only scope of turning f_n into a *real* polynomial in the case d is odd). We know from Theorem 28 that for the generic choice of $W \subset \mathfrak{so}(d)$ and $p \in \Delta^2$ each Lagrange multiplier belongs to some $Z(f_{k_1}, \ldots, f_{k_\nu})$, where ν is the number of positive integer eigenvalues of the matrix $i\omega A$. Now if $u = u_1 + \cdots + u_\nu$ is the control associated to a geodesic with final point p and energy less than s, Proposition 36 implies:

$$k_j \leq sc_p \quad \text{for } j = 1, \ldots, \nu$$

for a constant c_p depending only on W and p.

Thus the way to get all *possible* Lagrange multipliers associated to geodesics ending at p with energy less than s is by intersecting ν of the hypersurfaces $Z(f_k)$ for $k = 1, \ldots, \lfloor sc_p \rfloor$ and $\nu = 1, \ldots, l$ and considering the points in this intersection where the normal space contains p. There are

$$\binom{\lfloor sc_p \rfloor}{\nu} = O(s)^\nu$$

possible ways of choosing the hypersurfaces and the above argument implies each such choice can contribute by at most β_ν Lagrange multipliers.

In particular the set of Lagrange multipliers for p with energy less than s is bounded by:

$$\text{Card}\{\text{Lagrange multipliers } \omega \text{ such that } \omega(p) \leq s\} \leq \sum_{\nu=1}^{l} O(s)^{\nu} = O(s)^{l}.$$

To each critical manifold with energy less than s there corresponds one and only one Lagrange multiplier whose scalar product with p is less than s, hence the conclusion follows. $\square$

As a corollary we derive now Morse-Bott inequalities; as already stated the following bound will be improved in the next section.

Corollary 39 (Morse-Bott inequalities). *For the generic choice of $W \subset \mathfrak{so}(d)$ and of the point $p \in \Delta^2$ we have:*

$$b(\Omega_p^s) \leq \sum_{J(C_\omega) \leq s} b(C_\omega) = O(s)^l.$$

Proof. By Lemma 29 and Theorem 30 we know that each critical manifold is an intersection of l quadrics in $\mathbb{R}^{2d}$ (it is the preimage of p under $q|_{E(\omega)}$): in particular the possible homeomorphism types of such manifold are finite and there is a constant β such that $b(C_\omega) \leq \beta$ for every critical manifold C_ω.

Since the sum $\sum_{J(C) \leq s} b(C)$ contains at most $O(s)^l$ terms (by Theorem 38), the conclusion follows from Morse-Bott inequalities (see Appendix 2.2). $\square$

3.2.3 Asymptotic total Betti number

The aim of this section is to refine the bound obtained from Morse-Bott inequalitis for $b(\Omega_p^s)$ given in Corollary 39.

We start by proving some technical results; roughly they tell that all the information about the critical manifolds with bounded Energy and their index can be found in a finite dimensional approximation of the space of controls; moreover it is shown the way that the dimension of the finite dimensional approximation depends on the sublevel of the Energy.

Proposition 40. *For the generic choice of $W \in \mathfrak{so}(d)$ and the generic point $p \in \Delta^2$ we have: (a) the critical manifolds of J on Ω_p with Energy less than s coincide with the critical manifolds of J restricted to $\Omega_p \cap T^{\lfloor sc_p \rfloor}$; (b) their index is the same either if they are considered critical manifolds for J or if they are considered critical manifolds for J restricted to $\Omega_p \cap T^{\lfloor sc_p \rfloor}$.*

Proof. We have already proved that there exists a constant c_p such that all the critical points with energy less than s are contained in $\Omega_p \cap T^{\lfloor sc_p \rfloor}$ (Proposition 36); since the spaces T_k are orthogonal with respect to both the quadratic maps J and q, the critical points of $J|_{\Omega_p \cap T^{\lfloor sc_p \rfloor}}$ are given by the same equations as for critical points of $J|_{\Omega_p}$ using the Lagrange multipliers rule.

Let ω be a Lagrange multiplier for p with energy less than s. The Hessian of the energy J is $\langle \mathrm{Id} - \omega Q \cdot, \cdot \rangle$, so its eigenvalues are $1 - \frac{\alpha_i(\omega)}{k}$ with $k \in \mathbb{N}_0$ and $\alpha_i(\omega)$ eigenvalues of ωA as usual. Therefore we have negative eigenvalues for every integer k such that $k < \alpha_i(\omega)$ for at least one i; since $\alpha_i(\omega) \leq \|\omega\|$, it follows from Lemma 37 (as in the proof of Proposition 36) that

$$k \leq \|\omega\| \leq c_p \langle \omega, p \rangle \leq sc_p$$

with the same constant c_p as in Proposition 36. □

The next proposition tells also all the topological information is contained in such a finite dimensional approximation.

Proposition 41. *For a generic choice of $W \subset \mathfrak{so}(d)$ and of $p \in \Delta^2$ there exists a constant $r_p > 0$ such that for every $m \in \mathbb{N}$:*

$$\Omega_p^s \quad \textit{deformation retracts to} \quad \Omega_p^s \cap T^{\lfloor sr_p \rfloor + m}.$$

In particular: $H_*(\Omega_p^s) \simeq H_*(\Omega_p^s \cap T^{\lfloor sr_p \rfloor + m})$.

Proof. Given $L \in \mathbb{N}$ we can define the function f_L, "distance from T^L" in the following way: every $u \in \Omega_p^s$ can be uniquely written as $u = \bar{u} + v$ where $\bar{u} \in T^L$ and $v \in \left(T^L\right)^\perp$; then we define $f_L(u) \doteq \|v\|^2$. Assume that for a suitable L there are no critical points for f_L outside $\Omega_p^s \cap T^L$ and the function f_L satisfies the Palais-Smale condition: then we can retract the manifold $f_L^{-1}([0, s]) = \Omega_p^s$ on the sublevel set $f_L^{-1}(0) = \Omega_p^s \cap T^L$ by Theorem 4 (the function f_L is bounded on $\widehat{\Omega}_{\epsilon p}$ since $f_L \leq J \leq s$). The manifolds we are considering have boundary, but we can apply the argument above to $\Omega_p \cap \{J < s + \delta\}$: it deformation retracts to Ω_p^s again by Theorem 4 if we choose δ small enough not to have new critical values for the Energy; we can indeed choose such δ small enough not to have new critical values for f_L as well.

We are now going to prove that f_L satisfies the Palais-Smale condition and to find for which L we are sure not to have critical points for f_L in $\Omega_p^s \setminus (\Omega_p^s \cap T^L)$.

The gradient of the function f_L at $u = \bar{u} + v$ is $2v$ restricted to $T_u\Omega_p$: this means that there exists $\eta \in (\Delta^2)^*$ such that $\nabla_u F_L = 2v - 2\eta Q u$. For every critical point u for f_L there exists η such that $v-\eta Qu = v-\eta Qv-\eta Q\bar{u} = 0$: since the space T^L and its orthogonal are invariant with respect to ηQ we have the equivalent couple of conditions

$$v = \eta Qv, \quad \eta Q\bar{u} = 0;$$

the first condition tells that v is a geodesic going somewhere, the second one tells that $q(\bar{u})$ is a critical value for q.

We need f_L to satisfy the Palais-Smale condition: having the explicit expression of its gradient ∇f_L, we omit this verification whose proof is analogous to the one for J in Theorem 31.

We are going to prove that there exists a constant r_p such that if we take $L = \lfloor sr_p \rfloor$ there are no critical points for f_L in $\Omega_p^s \setminus (\Omega_p^s \cap T^L)$.

Assume that such a constant does not exist: then for every term ρ_n of a diverging sequence of positive real numbers, we find $s_n > 0$ and a critical point $u(n)$ for the function $f_{\lfloor s_n\rho_n \rfloor}$ outside $T^{\lfloor \rho_n s_n \rfloor}$ with Energy less than s_n. By hypothesis $v(n) \in \left(T^{\lfloor \rho_n s n_n \rfloor}\right)^{\perp}$ so that its Fourier expansion is $v(n) = \sum_{k>\rho_n s_n} v_k(n)$. Recall that $P \doteq Q|_{T_1}$: then we have:

$$\|q(v(n))\| = \left\| \sum_{k>\rho_n s_n} \frac{1}{k}\langle Pv_k(n), v_k(n)\rangle \right\| \leq \sum_{k>\rho_n s_n} \frac{1}{k}\|P\|\|u_k(n)\|^2 \leq$$

$$\leq \frac{1}{\rho_n s_n}\|P\|\|v(n)\|^2 \leq \frac{1}{\rho_n s_n}\|P\| s_n = \frac{\|P\|}{\rho_n}.$$

The last term of the chain of inequalities converges to zero: it follows that $q(v(n))$ goes to zero as well. Since $p = q(\bar{u}(n)) + q(v(n))$, we get

$$\lim_{n\to\infty} q(\bar{u}(n)) = p,$$

and as we noticed before $\bar{u}(n)$ is a critical point for q. This means that we have a sequence of critical values for q converging to p, which is impossible since we picked our p into an open subset of regular values (i.e. $p \in \Delta^2\backslash\Sigma_1$). We end the proof by noticing that the condition for having critical points for $f_{\lfloor r_p s_n \rfloor}$ has to be verified on every T_k separately (they are orthogonal and invariant with respect to Q): since there are no critical points of $f_{\lfloor r_p s_n \rfloor}$ outside $T^{\lfloor r_p s_n \rfloor}$, there are no critical points of the same function restricted to $T^{\lfloor r_p s_n \rfloor + m}$ (similarly to what happens to the critical points of J in Proposition 40) and it follows eventually that $\Omega_p^s \cap T^{\lfloor r_p s_n \rfloor + m}$ deformation retracts onto $\Omega_p^s \cap T^{\lfloor r_p s_n \rfloor}$ for every $m \in \mathbb{N}$. $\square$

As a corollary we get the following interesting result, that controls the growth rate of the index of the highest nonzero Betti number of Ω_p^s.

Corollary 42. *For a generic $W \subset \mathfrak{so}(d)$ and $p \in \Delta^2$, there exists a constant $r_p > 0$ such that:*

$$\max_i \{i \,|\, b_i(\Omega_p^s) \neq 0\} \leq 2d \lfloor r_p s \rfloor .$$

Proof. By Proposition 41 there exists $c_p > 0$ such that $H_*(\Omega_p^s) \simeq H_*(\Omega_p^s \cap T^{\lfloor c_p s \rfloor})$. In particular Ω_p^s has the homology of a semialgebraic subset of $\mathbb{R}^{2d\lfloor c_p s \rfloor}$ (namely $\Omega_p^s \cap T^{\lfloor c_p s \rfloor}$) and its j-th Betti number must be zero for $j > 2d\lfloor c_p s \rfloor$. □

Everything is now ready for the proof of the main theorem of this section.

Theorem 43 (Sharper bound on total Betti number)**.** *For the generic choice of $W \subset \mathfrak{so}(d)$ and of $p \in \Delta^2$ we have:*

$$b(\Omega_p^s) = O(s)^{l-1}.$$

Proof. First we know from Proposition 41 that there exists $r_p > 0$ such that:

$$b(\Omega_p^s) = b(\Omega_p^s \cap T^{\lfloor r_p s \rfloor}) \quad \text{for all } s > 0 \text{ and } m \in \mathbb{N}.$$

It means that Ω_p^s has the same Betti numbers as

$$\{x \in T^{\lfloor r_p s \rfloor} \,|\, q_1(x) = p_1, \ldots, q_l(x) = p_l\} \cap \{\|x\|^2 \leq 2s\}.$$

Notice that the dimension of $T^{\lfloor r_p s \rfloor}$ is a $O(s)$.

Let us consider the semialgebraic set (a level set of a quadratic map with l components):

$$X = \{x \in T^{\lfloor r_p s \rfloor} \,|\, q_1(x) = p_1, \ldots, q_l(x) = p_l\}$$

and $\epsilon > 0$ small enough such that $X \cap \{2s - \epsilon \leq \|x\|^2 \leq 2s + \epsilon\}$ deformation retracts onto X (the existence of such ϵ is guaranteed by semialgebraic triviality, see [8]). Let also $A = X \cap \{\|x\|^2 \leq 2s + \epsilon$ (which deformation retracts onto $X \cap \{\|x\|^2 \leq 2s\}$ and $B = X \cap \{2s - \epsilon \leq \|x\|^2\}$. The Mayer-Vietoris exact sequence of the pair (A, B) gives $b(A) + b(B) \leq b(A \cap B) + b(A \cup B)$, which implies:

$$b(\Omega_p^s) = b(A) \leq b(A \cap B) + b(A \cup B).$$

Since $A \cap B$ deformation retracts onto $X \cap \{\|x\|^2 = 2s\}$, then it is defined by l quadratic equations on a sphere of dimension $\dim(T^{\lfloor r_p s \rfloor}) - 1 = O(s)$; on the other hand $X = A \cup B$ is given by l quadratic equations in a vector space of dimension $\dim(T^{\lfloor r_p s \rfloor}) = O(s)$; hence by Proposition 13 the total Betti numbers of these spaces are bounded by:

$$b(A \cap B) = O(s)^{l-1} \quad \text{and} \quad b(A \cup B) = O(s)^{l-1}$$

and the conclusion follows. □

3.2.4 A topological coarea formula

As we found in the previous section, we have an upper bound for the total Betti number of the sublevels of the Energy; what we do not know is wether this bound gives the exact growth rate of $b(\widehat{\Omega}^s_{\epsilon p})$ or it is still overcounting, as Morse-Bott inequalities do. In this section we compute the first order asymptotic of $b(\Omega^s_p)$ in s, for the case $l = 2$. It turns out that for a generic choice of the Carnot group structure $W \subset \mathfrak{so}(d)$ and the point $p \in \Delta^2$, the leading term is a *real* number and can be analytically computed using only the data $W = \text{span}\{A_1, A_2\}$. It follows that the previous bound is sharp, at least in the case of codimension $l = 2$.

Consider a unit circle $S^1 \subset W$. For a generic W the eigenvalues of ωA are distinct and differentiable almost everywhere (the set of matrices in $\mathfrak{so}(d)$ with multiple eigenvalues is a cone with codimension 3, and the eigenvalues are semialgebraic functions of the parameter $\omega \in S^1$). Thus there exist semialgebraic functions $\alpha_j : S^1 \to \mathbb{R}$ such that the $\alpha_j(\omega)$, for $j = 1, \ldots, d$, are the coefficients of the canonical form of ωA. Given $p \in \Delta$ we consider the *rational* functions $\lambda_j : S^1 \to \mathbb{R} \cup \{\infty\}$ given by:

$$\lambda_j : \omega \mapsto \left| \frac{\alpha_j(\omega)}{\langle \omega, p \rangle} \right| \quad \text{for} \quad j = 1, \ldots, d.$$

Notice that when ω approaches $p^\perp$ these functions might explode, that is why they are rational in ω; on the other hand they are semialgebraic and differentiable almost everywhere and it makes sense to consider the integral:

$$\tau(p) \doteq \frac{1}{2} \int_{S^1} \sum_{j=1}^{d} \left| \dot{\lambda}_j(\omega) \right| - \left| \sum_{j=1}^{d} \dot{\lambda}_j(\omega) \right| d\omega. \tag{3.10}$$

The convergence of the integral follows from the fact that where the derivatives of the λ_js explode, they all have the same sign and the integrand vanishes.

The next theorem proves that for a fixed p, as a function of s:

$$b(\Omega^s_p) = \tau(p)s + o(s) \quad \text{as } s \to \infty.$$

Theorem 44. *If the corank $l = 2$, for a generic choice of $W \subset \mathfrak{so}(d)$ and $p \in \Delta^2$ we have:*

$$\lim_{s \to \infty} \frac{b(\Omega^s_p)}{s} = \tau(p).$$

Proof. In order to compute the asymptotic for $b(\Omega_p^s)$ for $s \to \infty$ we use (2.5). In fact we have seen that Ω_p^s is homotopy equivalent to $\{v \in H \,|\, q(v) = p/s,\ \|v\|^2 = 1\}$ and the latter can be rewritten as:

$$\left\{v \in H \,|\, q(v) - \frac{\|v\|^2}{s} p = 0\right\} \cap S$$

where S is the infinite dimensional sphere $S = \{\|v\|^2 = 1\}$. In particular we can present our set as the intersection of two quadrics in H on the unit sphere S; thus we let i_s^- for the index function of the quadratic map $q - p\|\cdot\|^2/s$ (we are using the notation of 14). In this setting the set $P \subset S^1$ coincides with $\{\omega \in S^1 \,|\, \langle \omega, p\rangle < 0\}$. In fact if we let $p = (p_1, p_2)$, here the two quadrics we are considering are $q_1 - p_1\|\cdot\|^2/s$ and $q_2 - p_2\|\cdot\|^2/s$ and for every ω the self-adjoint operator on H corresponding to the quadratic form ωq is $\omega Q - \omega(p)\mathbb{1}/s$. In particular the spectrum of $\omega Q - \omega(p)\mathbb{1}/s$ is obtained by translating the spectrum of ωQ by $\langle \omega, p\rangle/s$ and since ωQ is compact and its spectrum is symmetric with respect to the origin, we see that in order to have finitely many negative eigenvalues we need $\langle \omega, p\rangle < 0$.

On the other hand the subspaces T_k are invariant by both ωQ and any multiple of the identity, thus the index function can be computed as:

$$\begin{aligned}\mathrm{i}_s^-(\omega) &= \sum_{k\geq 1} \mathrm{i}^-\left(\omega Q - \frac{\langle \omega, p\rangle}{s}\mathbb{1}|_{T_k}\right) = \sum_{k\geq 1} \mathrm{i}^-\left(\frac{\omega Q_0}{k} - \frac{\langle \omega, p\rangle}{s}\mathbb{1}\right) = \\ &= \sum_{j=1}^{d} \left\lfloor \frac{s\alpha_j(\omega)}{\langle \omega, p\rangle} \right\rfloor = \sum_{j=1}^{d} \lfloor s\lambda_j(\omega)\rfloor\end{aligned}$$

where in the second line we have used the fact that the spectrum of $\frac{\omega Q_0}{k} - \frac{\langle \omega,p\rangle}{s}\mathbb{1}$ is of the form $\frac{\alpha_j(\omega)}{k} - \frac{\langle \omega,p\rangle}{s}$. In the sequel we also identify $P \subset S^1$ with a subset of $[0, 2\pi]$ in the standard way.

Denoting now by $\mu(s)$ the number of local maxima of i_s^- on P, we see that formula (2.5) implies:

$$b(\Omega_p^s) = 2\mu(s) + 1 - b_0(P_0) = 2\mu(s) + O(1) \tag{3.11}$$

In fact, using the long exact sequence of the pair (P_{j+1}, P_j), we can rewrite $b_0(P_{j+1}, P_j) = b_0(P_{j+1}) - b_0(P_j) + b_1(P_{j+1}, P_j)$; substituting these identities into $b(\hat{\Omega}_{\epsilon p}) = 1 + \sum_{j\geq 1} b_0(P_{j+1}, P_j) + b_1(P_{j+2}, P_{j+1})$ we get $b(\hat{\Omega}_{\epsilon p}) = 1 - b_0(P_0) + 2\sum_{j\geq 1} b_1(P_{j+1}, P_j)$. Since each local maximum of i_ϵ^- contributes by 1 to one of the $b_1(P_{j+1}, P_j)$ and $b(P_0) \leq 1$ (since P_0 is convex), then (3.11) follows.

In order to compute the asymptotic of the number of maxima of i_ϵ^- we introduce the following auxiliary data. First we let $\lambda = \sum_{j=1}^d \lambda_j(\omega)$ and notice that this is a semialgebraic function. In particular we can divide P into a finite number of intervals (arcs):

$$P = (\omega_0, \omega_1] \cup [\omega_1, \omega_2] \cup \cdots \cup [\omega_m, \omega_{m+1}] \cup [\omega_{m+1}, \omega_{m+2})$$

such that for every j, k the functions α_j as well as α are monotone on (ω_k, ω_{k+1}). Labeling $I_k = [\omega_k, \omega_{k+1}]$ we see that also each $\lfloor s\lambda_j \rfloor$ is monotone on I_k. On the other hand monotonicity of i_s^- is granted only where the signs of the derivatives of the λ_j all agree. Since for the generic choice of p the functions α_j do not vanish on $\{\omega_0, \omega_{m+2}\}$ (the orthogonal complement of p on S^1), λ_j approaches infinity when approaching ω_0 or ω_{m+2}; in particular i_s^- *is* monotone on I_0 and I_{m+1} and has no local maxima on them.

For every $j \in \{1, \ldots, m\}$ let us denote respectively by $\mu_j(s)$ and $\sigma_j(s)$ the number of local maxima of i_s^- on I_j and the number of subintervals of I_j where i_s^- is constant (thus $\sigma_j(s)$ equals the number of "jumps" of the integer valued function i_s^- on I_j).

For every interval $I_j = [\omega_j, \omega_{j+1}]$ we see that:

$$|\mathrm{i}_s^-(\omega_{j+1}) - \mathrm{i}_s^-(\omega_j)| = \sigma_j(s) - 2\mu_j(s)$$

In particular summing all these equations and using the fact that i_s^- is monotone on I_0 and I_{m+1}, combining with (3.11) we get:

$$\begin{aligned}\frac{b(\Omega_p^s)}{s} &= 2\frac{\mu(\epsilon)}{s} + O(1/s) = 2\sum_{j=1}^m \frac{\mu_j(s)}{s} + O(1/s) = \\ &= \sum_{j=1}^m \frac{\sigma_j(s)}{s} - \sum_{j=1}^m \left| \frac{\mathrm{i}_s^-(\omega_{j+1}) - \mathrm{i}_s^-(\omega_j)}{s} \right| + O(1/s).\end{aligned}$$

Now we notice that as $s \to \infty$, the function i_s^-/s converges uniformly to λ, thus:

$$\lim_{s\to\infty} \sum_{j=1}^m \left| \frac{\mathrm{i}_s^-(\omega_{j+1}) - \mathrm{i}_s^-(\omega_j)}{s} \right| + O(1/s) = \sum_{j=1}^m |\lambda(\omega_{j+1}) - \lambda(\omega_j)| = \tag{3.12}$$

$$= \int_{\omega_1}^{\omega_{m+1}} \left| \dot{\lambda}(\omega) \right| d\omega. \tag{3.13}$$

It remains to evaluate $\lim_s \sum_j \frac{\sigma_j(s)}{s}$. To this end we let $\sigma_j^i(s)$ be the number of jumps of $\lfloor s\lambda_i / \rfloor$ on the interval I_j. We notice that $\sigma_j^i(s) =$

$\sum_{i=1}^{d} \sigma_j^i(s) + O(1)$: in fact i_s^- jumps exactly when one of the $\lfloor s\lambda_j \rfloor$ jumps and these function all jump at different points (except for the points where two eigenvalues are in resonance, but these are in finite number bounded independently of s); we also notice that each function $\lfloor s\lambda_i \rfloor / s$ converges uniformly to λ_i. Thus we get:

$$\lim_{s\to\infty} \sum_{j=1}^{m} \frac{\sigma_j(s)}{s} = \lim_{s\to\infty} \sum_{j=1}^{m} \left(\sum_{i=1}^{d} \frac{\sigma_j^i(s)}{s} \right) = \tag{3.14}$$

$$= \lim_{s\to\infty} \sum_{j=1}^{m} \left(\sum_{i=1}^{d} \left| \frac{\lfloor s\lambda_i(\omega_{j+1}) \rfloor}{s} - \frac{\lfloor s\lambda_i(\omega_j) \rfloor}{s} \right| \right) = \tag{3.15}$$

$$= \sum_{j=1}^{m} \left(\sum_{i=1}^{d} |\lambda_i(\omega_{j+1}) - \lambda_i(\omega_j)| \right) =$$

$$= \sum_{j=1}^{m} \int_{\omega_j}^{\omega_{j+1}} \left(\sum_{i=1}^{d} \left| \dot{\lambda}_i(\omega) \right| \right) d\omega = \tag{3.16}$$

$$= \int_{\omega_1}^{\omega_{m+1}} \left(\sum_{i=1}^{d} \left| \dot{\lambda}_i(\omega) \right| \right) d\omega. \tag{3.17}$$

Combining (3.12) and (3.14) we finally get:

$$\lim_{s\to\infty} \frac{b(\Omega_p^s)}{s} = \int_{\omega_1}^{\omega_{m+1}} \sum_{i=1}^{d} \left| \dot{\lambda}_i(\omega) \right| - \left| \sum_{i=1}^{d} \dot{\lambda}_i(\omega) \right| d\omega$$
$$= \int_{P} \sum_{i=1}^{d} \left| \dot{\lambda}_i(\omega) \right| - \left| \sum_{i=1}^{d} \dot{\lambda}_i(\omega) \right| d\omega$$

where the last identity follows from the fact that on I_0 and I_{m+1} the two functions $\sum_{i=1}^{d} \left| \dot{\lambda}_i(\omega) \right|$ and $\left| \sum_{i=1}^{d} \dot{\lambda}_i(\omega) \right|$ are equal. The limit of the statement simply follows by noticing that $\alpha_i(\omega) = \alpha_i(-\omega)$ (i. e. the positive eigenvalues of $i\omega A$ are π-periodic). □

As a corollary we get the following result: it says that the topology of the set of paths reaching the point ϵp with energy $J \leq 1$ explodes; in particular the number of geodesics gets unbounded as well.

Corollary 45. *For a generic choice of $W \subset \mathfrak{so}(d)$ and the point $p \in \Delta^2$:*

$$\lim_{\epsilon\to 0} b(\Omega_{\epsilon p} \cap \{J \leq 1\}) = \infty.$$

Proof. First notice that $\Omega_{\epsilon p} \cap \{J \leq 1\}$ is homeomorphic to $\Omega_p \cap \{J \leq 1/\epsilon\}$; thus we can apply the above theorem.

In order to prove the limit it is enough to show that for the generic choice of W and p the integral $\tau(p)$ is not zero. Since the integrand function is always non-negative, it's enough to prove it doesn't vanish identically. Pick two distinct eigenvalues (functions), say $i\alpha_1$ and $i\alpha_2$, for the family $\{\omega A\}_{\omega \in S^1}$. Since these functions are continuous semialgebraic, $i\alpha_1$ has a maximum point $\overline{\omega}$ and we can assume this is not a critical point for α_2 also (this is a generic condition). Then in a neighborhood of $\overline{\omega}$ the derivatives of the corresponding λ_1 and λ_2 have different signs and the integrand is nonzero. $\square$

We conclude the section with an example where the topological coarea formula can be computed directly.

Example 1 (Commuting matrices, corank $l = 2$). Let us fix the corank $l = 2$. If the matrices A_1 and A_2 commute, they can be written simultaneously in their canonical form

$$A_i = \operatorname{diag}\left(v_1^i J_2, \ldots, v_k^i, 0_h\right),$$

where as usual J_2 is the 2×2 sympletic matrix and 0_h is the $h \times h$ zero matrix (possibly with $h = 0$. Setting $v_j \doteq (v_j^1, v_j^2)$ and given $\omega \in \mathbb{R}^{2*}$, the eigenvalues of the matrix ωA are $\pm\langle \omega, v_j \rangle$. Now we pick a generic $p \in \Delta^2$: having parametrized by t the unit circle in R^{2*}, the functions we need in order to compute $\tau(p)$ are $\lambda_j(t) \doteq \left| \frac{\langle \omega(t), v_j \rangle}{\langle \omega(t), p \rangle} \right|$, their derivatives being

$$\dot{\lambda}_j(t) = \frac{\operatorname{sgn}\left(\langle \omega(t), v_j \rangle\right)}{\operatorname{sgn}(\langle \omega(t), p \rangle)} \cdot \frac{\langle \dot{\omega}(t), v_j \rangle \langle \omega(t), p \rangle - \langle \omega(t), v_j \rangle \langle \dot{\omega}(t), p \rangle}{\langle \omega(t), p \rangle^2}.$$

Since the curve $\omega(t)$ is the arc-length parametrization of the unit circle, for every t the covectors $\omega(t), \dot{\omega}(t)$ form an orthonormal basis for $(\Delta^2)^*$; it follows that the term

$$m_j \doteq \langle \dot{\omega}(t), v_j \rangle \langle \omega(t), p \rangle - \langle \omega(t), v_j \rangle \langle \dot{\omega}(t), p \rangle$$

is the determinant of the matrix $\begin{pmatrix} p^1 & a_j^1 \\ p^2 & a_j^2 \end{pmatrix}$ which does not depend on t. Since the functions λ_j are periodic by π, we can modify the formula (3.10) by integrating on the semicircle $\{\langle \omega(t), p \rangle > 0\}$ and eliminating the coefficient of $1/2$ before the integral sign.

These two remarks allow us to simplify the expression of $\dot{\lambda}(t)$:

$$\dot{\lambda}_j(t) = \operatorname{sgn}(\langle \omega(t), v_j \rangle) \cdot \frac{m_j}{\langle \omega(t), p \rangle^2}.$$

Now, in order to have the term $\tau(p) > 0$, the integrand of

$$\begin{aligned}\tau(p) &= \int_{\langle\omega(t),p\rangle>0} \sum_{j=1}^{d} \left|\dot{\lambda}_j(\omega)\right| - \left|\sum_{j=1}^{d} \dot{\lambda}_j(\omega)\right| d\omega = \\ &= \int_{\langle\omega(t),p\rangle>0} \frac{1}{\langle\omega(t),p\rangle^2} \cdot \left(\sum_{i=1}^{d} |m_j| - \left|\sum_{i=1}^{d} \operatorname{sgn}(\langle\omega(t),v_j\rangle)\ m_j\right|\right) d\omega\end{aligned}$$

must be strictly positive somewhere. This happens if and only if there exist i and j such that for some t_0

$$\operatorname{sgn}(\langle\omega(t_0),v_i\rangle)m_i \cdot \operatorname{sgn}(\langle\omega(t_0),v_j\rangle)m_j,$$

namely the two factors have opposite sign. If $m_i \cdot m_j > 0$ we must have t_0 such that

$$\operatorname{sgn}(\langle\omega(t),v_i\rangle) \cdot \operatorname{sgn}(\langle\omega(t),v_j\rangle) < 0 :$$

if such a t_0 does not exist, we have

$$\operatorname{sgn}(\langle\omega(t),v_i\rangle) \cdot \operatorname{sgn}(\langle\omega(t),v_j\rangle) > 0$$

for all t such that $\langle\omega(t),p\rangle > 0$, meaning that v_i is proportional to v_j; a similar argument holds if $m_i \cdot m_j < 0$.

In order to have $\tau(p) = 0$ the only possibility is that every v_i is proportional to every other v_j, which is equivalent to say that A_1 and A_2 are proportional: this implies that $\dim\Delta^2 - \dim\Delta = 1$ which contradicts the hypothesis.

3.3 General sub-Riemannian manifolds

3.3.1 Homotopy equivalence of path spaces

Since Carnot groups approximate general sub-Riemannian structures, we expect that some properties of the topology of horizontal path spaces persist when we consider small enough neighborhood of a fixed point p; also we have to look at small enough neighborhood of the constant path in the space of horizontal paths.

Since we will study properties that depend on the germ of the distribution Δ at a regular point $p \in M$, we can assume without loss of generality that the manifold is $\mathbb{R}^n$ and the starting point is the origin; we have to assume that the origin is a regular point of the distribution Δ. In this case we consider privileged coordinates in a neighborhood U of the origin and consequently a 1-parameter family of dilations δ_ϵ and a nilpotent approximation; let us denote f the End-point map and $\hat{f}$ the End-point map associated to the nilpotent approximation. In this section we make the choice to denote with the $\hat{}$ the objects in the nilpotent approximation; for instance while Δ will be the general distribution, we will denote $\widehat{\Delta}$ the distribution spanned by the homogeneous parts of the vector fields that generate Δ, i.e. the nilpotent approximation.

Given the family of End-point maps $f_\epsilon(u) \doteq \delta_{1/\epsilon}(\epsilon u)$, we have the following Theorem [4]:

Theorem 46. *The maps f_ϵ converge in C^∞ topology to $\hat{f}$ when $\epsilon \to 0$.*

Since the End-point map f can be deformed smoothly to the End-point map $\hat{f}$ of the chosen nilpotent approximation, we expect when the parameter ϵ goes to 0, their fibers (the horizontal path spaces) have something in common. Indeed, we prove the existence of a weak homotopy equivalence between bounded energy horizontal path spaces of the two structure, for a suitable choice of the ending point.

On one hand, by acting with the dilations we prove the following

Lemma 47. *The following diffeomorphism holds:*

$$f_\epsilon^{-1}(p) \cap \{J \leq s\} \cong f^{-1}(\delta_\epsilon(p)) \cap \{J \leq \epsilon^2 s\}$$

Proof. Let us consider the following chain of equalities:

$$\begin{aligned} f_\epsilon^{-1}(p) \cap \{J \leq s\} &= \left\{ u \;\mid\; \frac{1}{2}\|u\|^2 \leq s,\ \delta_{1/\epsilon} f(\epsilon u) = p \right\} = \\ &= \left\{ u \;\mid\; \frac{1}{2}\|u\|^2 \leq s,\ f(\epsilon u) = \delta_\epsilon(p) \right\} = \\ &= \left\{ \frac{1}{\epsilon} v \;\mid\; \frac{1}{2}\|v\|^2 \leq \epsilon^2 s,\ f(v) = \delta_\epsilon(p) \right\} \cong \\ &\cong \left\{ v \;\mid\; \frac{1}{2}\|v\|^2 \leq \epsilon^2 s,\ f(v) = \delta_\epsilon(p) \right\} = \\ &= f^{-1}(\delta_\epsilon(p)) \cap \{J \leq \epsilon^2 s\} \end{aligned}$$

where the homeomorphism between the expressions in the third and fourth lines follows from a dilation of parameter ϵ in the space of controls. □

On the other hand the following Lemma relates the topology of the $f_\epsilon^{-1}(p)$ to the topology of $\hat{f}^{-1}(p)$:

Lemma 48. *For any given $s > 0$ there exists ϵ_s such that for every $\epsilon < \epsilon_s$ the following weak homotopy equivalence holds:*

$$f_\epsilon^{-1}(p) \cap \{J \le s\} \sim \hat{f}^{-1}(p) \cap \{J \le s\}$$

Proof. We define the following map:

$$\begin{aligned} F :& (0,1] \times \mathcal{H} \longrightarrow M \\ &(\epsilon, u) \mapsto f_\epsilon(u) \end{aligned}$$

Thanks to the C^∞ convergence the map F can be extended for $\epsilon = 0$ by letting $F(0,u) = \hat{f}(u)$.

Now we fix a vertical[2] point $p \in M$ which is a regular value for $\hat{f}$. It follows that p is also regular for the map F, after restricting his domain to $[0,T]$ for a suitable $T \le 1$; we want to prove that the homology of $\hat{f}^{-1}(p) = f_0^{-1}(p)$ doesn't change for small values of ϵ.

To this end we consider the projection of ϵ restricted to the fibers of F:

$$\begin{aligned} \pi_\epsilon : F^{-1}(p) &\longrightarrow [0,1] \\ (\epsilon, u) &\mapsto \epsilon. \end{aligned}$$

We study the critical values of π_ϵ by means of the Lagrange multiplier rule: (ϵ_0, u_0) is a critical point for π_ϵ iff there is a $\lambda \in T_p^*M$ such that

$$d_{(\epsilon_0,u_0)}\pi_\epsilon - \lambda \circ D_{(\epsilon_0,u_0)}F = 0; \tag{3.18}$$

The differential of π_ϵ is just the identity on the ϵ coordinate, while for the differential of F we have:

$$D_{(\epsilon_0,u_0)}F(\zeta, v) = \frac{\partial F}{\partial \epsilon}_{(\epsilon_0,u_0)}\zeta + D_{u_0}f_{\epsilon_0}(v).$$

The differential in equation (3.18) must be 0 for every vector (ζ, v) and in particular for the vectors of the kind $(0,v)$; it follows that $\lambda \circ D_{u_0}f_{\epsilon_0} = 0$, meaning that f_{ϵ_0} has a critical point mapped to p.

By hypotesis we know that p is not a critical value of $\hat{f}$. For every $u \in \hat{f}^{-1}(p) \cap \{J \le s\}$ we can find a finite dimensional subspace $L_u \subset \mathcal{H}$ such that the restriction of the differential of $\hat{f}$ is already surjective. Thanks to the continuity of the differential $D_u\hat{f}$ as a function of u in the weak topology, we can find a neighbourhood $\mathcal{O}_u$ of u such that the differential of $\hat{f}$ is surjective when restricted to the same linear subspace L_u. Since the ball $\{J \le u\}$ is

[2]Vertical with respect to the given nilpotent approximation.

compact in the weak topology, we can extract a finite subcover $\mathcal{O}_{u_1}, \dots, \mathcal{O}_{u_m}$: it follows that the differential of $\hat{f}$ is globally surjective on $\hat{f}^{-1}(p)$ if restricted to $L_0 \doteq \bigcup_{i=1}^m L_{u_i}$. Moreover, thanks to the finite dimension and the continuity of the differential Df_ϵ with respect to ϵ, we can find ϵ_s such that also the differential $D_u f_\epsilon$ restricted to L_0 is surjective for every $\epsilon \in [0, \epsilon_s)$ and every $u \in f_\epsilon^{-1}(p) \cap \{J \leq s\}$.

Let us consider the following direct system of finite dimensional linear subspaces of $\mathcal{H}$: $\mathcal{F} = \{L \subset \mathcal{H}, L_0 \subset L\}$, which is clearly cofinal: for every $L \in \mathcal{H}$, the manifolds $f_\epsilon^{-1} \cap \{J < s\} \cap L$ are diffeomorphic (the restriction of π_ϵ is a fibration) and in particular they have the same homotopy type, hence

$$\pi_*(f_\epsilon^{-1}(p) \cap L) = \pi_*(\hat{f}^{-1}(p) \cap L).$$

Since the homotopy groups of the spaces $f_\epsilon^{-1}(p) \cap \{J \leq s\}$ can be retrieved by taking the direct limit of the homotopy groups of a cofinal direct system of finite dimensional subspaces, we have

$$\pi_*(f_\epsilon^{-1}(p)) = \varinjlim \pi_*(f_\epsilon^{-1}(p) \cap L) = \varinjlim \pi_*(\hat{f}^{-1}(p) \cap L) = \pi_*(\hat{f}^{-1}(p)).$$

□

Lemma 47 and Lemma 48 give the following:

Theorem 49. *Given a generic rank d distribution Δ on $\mathbb{R}^n$ and privileged coordinates (x^α, y^i) together with the induced family of dilations δ_ϵ, if p is a regular value of the End-point map $\hat{f}$ associated to the nilpotent approximation the following homotopy equivalence holds:*

$$f^{-1}(\delta_\epsilon(p)) \cap \{J \leq \epsilon^2 s\} \sim \hat{f}^{-1}(p) \cap \{J \leq s\}.$$

Now that we have a way to relate the topology of horizontal paths of sub-Riemannian manifolds with their nilpotent approximation's counterpart with the aid of the 1-parameter family of dilations, we have to know which curves based on the origin can be realized as its orbits; moreover we have to find out whether there exists a choice of privileged coordinates such that the orbit is made of vertical points, in order to use our previous results. Fortunately, we have the following proposition:

Proposition 50. *For any parametric curve $\epsilon \mapsto \gamma(\epsilon)$ such that $\gamma(0) = 0$ and $\dot{\gamma}(0) \notin \Delta$, there exist privileged coordinates (x^α, y^i) on a neighbourhood U of 0 such that the curve γ in U is the coordinate curve y^l.*

Proof. The parametric curve restricted to a neighbourhood U of 0 is a submanifold of $\mathbb{R}^n$, so by the implicit function theorem we can find coordinates (x^α, y^i) such that the trajectory of the curve is given by $x^\alpha = 0, y^i = 0$ for $\alpha = 1, \ldots, d$ and $i = 1, \ldots, l-1$. After a linear change of coordinates we can assume that $X_\alpha|_0 = \frac{\partial}{\partial x^\alpha}|_0$; this is possible also because by hypotesis $\frac{\partial}{\partial y^i}(0)$ does not belong to the distribution Δ. Since

$$dy^i(X_\alpha)|_0 = X_\alpha(y^i)|_0 = \frac{\partial y^i}{\partial x^\alpha} = 0,$$

the condition of being privileged coordinates is satisfied by (x^α, y^i). $\square$

Studying the horizontal path spaces we found that we had to exclude some directions belonging to a measure zero set. We then must give a condition for the existence of privileged coordinates to have a given curve γ be realized as a vertical curve outside the "bad ones".

Proposition 51. *Given privileged coordinates (x^α, y^i) around 0 and the induced nilpotent approximation $\hat{X}^\alpha$, the map*

$$\psi : \Delta_0 \oplus [\Delta, \Delta]_0/\Delta_0 \longrightarrow \hat{\mathfrak{g}}$$

defined as $X_\alpha|_0 \mapsto \hat{X}_\alpha$ and $\left[\frac{\partial}{\partial y^i}\Big|_0\right] \mapsto \frac{\partial}{\partial y^i}$ is a Lie algebra isomorphism.

Proof. We can assume that the privileged coordinates satisfy the condition $X_\alpha|_0 = \frac{\partial}{\partial x^\alpha}|_0$ and we write down the Taylor expansion of X_α:

$$X_\alpha = \frac{\partial}{\partial x^\alpha} + a^i_{\alpha\mu} x^\mu \frac{\partial}{\partial y^i} + a^\nu_{\alpha\mu} x^\mu \frac{\partial}{\partial x^\nu} + a^i_{\alpha\mu\nu} x^\mu x^\nu \frac{\partial}{\partial y^i} + a^i_{\alpha j} y^j \frac{\partial}{\partial y^i} + R$$

where the term R has order ≥ 1; the homogeneous of order -1 vector fields are $\hat{X}^\alpha = \frac{\partial}{\partial x^\alpha} + a^i_{\alpha\mu} x^\mu \frac{\partial}{\partial y^i}$. We compute the brackets of the $\hat{X}^\alpha$ and we obtain:

$$[\hat{X}_\alpha, \hat{X}_\beta] = (a^i_{\alpha\beta} - a^i_{\beta\alpha}) \frac{\partial}{\partial y^i}.$$

Now we compute the brackets of the X^α in 0, so that we need only to consider only the terms of order -1 and 0:

$$[X^\alpha, X^\beta] = [\hat{X}_\alpha, \hat{X}_\beta] + (a^\nu_{\beta\alpha} - a^\nu_{\beta\alpha}) \frac{\partial}{\partial x^\nu} + R'$$

where the term R' vanishes at 0; so we have

$$[X^\alpha, X^\beta]\big|_0 = [\hat{X}_\alpha, \hat{X}_\beta]\Big|_0 = (a^i_{\alpha\beta} - a^i_{\beta\alpha}) \left.\frac{\partial}{\partial y^i}\right|_0 + (a^\nu_{\beta\alpha} - a^\nu_{\beta\alpha}) \left.\frac{\partial}{\partial x^\nu}\right|_0$$

which is equivalent to $(a^i_{\alpha\beta} - a^i_{\beta\alpha}) \left.\frac{\partial}{\partial y^i}\right|_0$ in Δ_0^2/Δ_0; this proves that ψ is a Lie algebra isomorphism. $\square$

We found that for a curve γ passing through the origin and transversal to the distribution Δ in the origin, the condition for the existence of privileged coordinates such that the curve is a "good" vertical coordinate curve is to have the projection under ψ of its tangent vector in the origin outside the set $\Sigma_1 \cup \Sigma_2 \cup \Sigma_3 \cup \Sigma_4$; since the latter has measure zero, it follows that the set of "bad" directions for the general sub-Riemannian structure has measure zero as well.

3.3.2 Low Energy geodesics

In this subsection we make use of the homotopy equivalence in the previous one in order to find a lower bound for the number of geodesics with bounded energy going to the point $\delta_\epsilon(p)$. From now on we set $p_\epsilon \doteq \delta_\epsilon(p)$.

Lemma 52. *If c_1 is a critical value for the Energy on the space of paths $\widehat{\Omega}_p$, all the values nc_1 for $n \in \mathbb{N}$ are critical.*

Proof. We have from Equation 3.6 at page 43 that controls associated to geodesics on a Carnot group going have the form

$$u(t) = e^{-t\omega A} v,$$

where $\omega \in (\Delta^2)^*$ is the Lagrange multiplier and $vin\Delta$. Given $n \in \mathbb{N}$, we define

$$u_n(t) \doteq e^{-nt\omega A} \sqrt{n}\, v,$$

and we look at the expression for its final point:

$$\begin{aligned}\hat{f}(u_n) &= \int_0^{2\pi} \left\langle \int_0^t e^{-n\tau\omega A} \sqrt{n}\, v \, d\tau, A e^{-nt\omega A} \sqrt{n}\, v \right\rangle dt = \\ &= \int_0^{2\pi} n \left\langle \int_0^t e^{-n\tau\omega A} v \, d\tau, A e^{-nt\omega A} v \right\rangle dt.\end{aligned}$$

By setting $t' = nt$ and $\tau' = n\tau$ we get

$$\begin{aligned}\hat{f}(u_n) &= \int_0^{2\pi n} \left\langle \frac{1}{n} \int_0^{t'} e^{-\tau'\omega A} \sqrt{n}\, v \, d\tau', A e^{-t'\omega A} \sqrt{n}\, v \right\rangle dt' = \\ &= \int_0^{2\pi} \left\langle \int_0^{t'} e^{-\tau'\omega A} \sqrt{n}\, v \, d\tau', A e^{-t'\omega A} \sqrt{n}\, v \right\rangle dt' = \hat{f}(u),\end{aligned}$$

where the last equality follows from the periodicity of $e^{-t'\omega A}$. From a single geodesic u with Lagrange multiplier ω we found a countable family of geodesics u_n with Lagrange multipliers $n\omega$ going to the same point; their Energy is given by their Lagrange multipliers as follows $J(u_n) = \left\langle n\omega, \hat{f}u_n \right\rangle = nJ(u)$, which concludes the proof. $\square$

In order to prove the main Theorem of this subsection we need a slight improvement of the Theorem 49, in order to deal with different sublevels of the Energy at the same time.

Lemma 53. *In the homotopy equivalence $\Omega_{p_\epsilon}^{\epsilon^2 s} \sim \widehat{\Omega}_p^s$ for $\epsilon < \epsilon_s$, the same ϵ_s works for smaller sublevels of the Energy: namely if $s' < s$,*

$$\Omega_{p_\epsilon}^{\epsilon^2 s'} \sim \widehat{\Omega}_p^{s'}$$

for every $\epsilon < \epsilon_s$; equivalently $\epsilon_{s'} \geq \epsilon_s$.

Proof. In the proof of Lemma 48 we found a subspace L_0 of the space of controls such that the restriction to L_0 of the differential $D_u f_\epsilon$ is surjective for every $\epsilon < \epsilon_s$ and every $u \in f_\epsilon^{-1}(p) \cap \{J \leq s\}$; then it is still true for every $u \in f_\epsilon^{-1}(p) \cap \{J \leq s'\}$ if $s' \leq s$, from which the conclusion follows. $\square$

In a few words we can say that the function $s \mapsto \epsilon_s$ is a non-increasing one.

Theorem 54. *Let Δ be a generic step 2 sub-Riemannian structure on $\mathbb{R}^n$ such that the origin 0 is a regular point; consider then a contractible open neighborhood $U \subset \mathbb{R}^n$ of 0. For every positive real number C and a generic point $p \in U$, the number of geodesics from x to p_ϵ with Energy less than C explodes as $\epsilon \to 0$, namely*

$$\lim_{\epsilon \to 0} \nu\left(\Omega_{p_\epsilon}^C\right) = \infty.$$

Proof. We know (Proposition 51) that for the generic p we find a nilpotent approximation $\hat{\Delta}$ of Δ such that p is vertical with respect to $\hat{\Delta}$. From Lemma 52 we know that if p is a vertical point in a Carnot group and c_1 is a critical value of the Energy, then all of its integer multiples $c_k \doteq kc_1$ are critical values. We choose two sequences of regular values s_k and S_k for the Energy such that the only critical value for the Energy in the interval $[s_k, S_k]$ is c_k. It follows that the relative homology $H(\tilde{\Omega}_p^{S_k}, \tilde{\Omega}_p^{s_k})$ is the same as the homology of the non-degenerate critical manifold associated to c_k and hence it is not trivial.

For any given $s > 0$, Theorem 49 gives the weak homotopy equivalence between $\Omega_{p_\lambda}^{\lambda^2 s}$ and $\widehat{\Omega}_p^s$. for every $\lambda < \epsilon_s$; from Lemma 53 it follows that for the same range of λ we have the following homotopy equivalences:

$$\Omega_{p_\lambda}^{\lambda^2 s_k} \sim \widehat{\Omega}_p^{s_k},\, \Omega_{p_\lambda}^{\lambda^2 S_k} \sim \widehat{\Omega}_p^{S_k}$$

for every $k \in \mathbb{N}$ such that $s_k < s$ and $S_k < s$. From the homotopy equivalence $\Omega_{p_\lambda}^{\lambda^2 s} \sim \widehat{\Omega}_p^s$ and the Five Lemma applied to the long exact sequence in homology induced by the inclusion $\widehat{\Omega}_p^{s_k} \hookrightarrow \widehat{\Omega}_p^{S_k}$, namely

$$\begin{array}{ccccccccc}
\xrightarrow{i_*} & H_i(\widehat{\Omega}_p^{S_k}) & \xrightarrow{p_*} & H_i(\widehat{\Omega}_p^{S_k}, \widehat{\Omega}_p^{s_k}) & \xrightarrow{\partial_*} & H_{i+1}(\widehat{\Omega}_p^{S_k}) & \xrightarrow{i_*} \\
& \downarrow & & \downarrow & & \downarrow & \\
\xrightarrow{i_*} & H_i(\Omega_{p_\lambda}^{\lambda^2 S_k}) & \xrightarrow{p_*} & H_i(\Omega_{p_\lambda}^{\lambda^2 S_k}, \Omega_{p_\lambda}^{\lambda^2 s_k}) & \xrightarrow{\partial_*} & H_{i+1}(\Omega_{p_\lambda}^{\lambda^2 s_k}) & \xrightarrow{i_*} & ,
\end{array}$$

we have the isomorphism

$$H_*\left(\widehat{\Omega}_p^{S_k}, \widehat{\Omega}_p^{s_k}\right) \cong H_*\left(\Omega_{p_\lambda}^{\lambda^2 S_k}, \Omega_{p_\lambda}^{\lambda^2 s_k}\right).$$

As a consequence we have that in the interval $[\lambda^2 s_k, \lambda^2 S_k]$ there must be a critical value for the Energy defined on Ω_{p_λ}, since the relative homology of the sublevels is non-zero. The numbers of these critical values is the integer k_s such $S_k \le s$ but $S_{k+1} > s$, which is $k_s = \lfloor s/s_1 \rfloor$.

Up to now we found that the number of geodesics going to p_λ with Energy less than $\lambda^2 s$ is k_s, that is $\nu\left(\Omega_{p_\lambda}^{\lambda^2 s}\right) \ge \lfloor s/s_1 \rfloor$. Given any real number $C > 0$, nothing changes if we apply to all the data the dilation $\delta_{C/s}$. On the space of controls associated to the nilpotent approximation it is the same to multiply the controls by $\sqrt{C/s}$: the space $\widehat{\Omega}_p^s$ becomes $\widehat{\Omega}_{p_{C/s}}^C$. By arguing as before we find that in the space $\Omega_{p_{\lambda C/s}}^{\lambda^2 C}$ there are at least $\lfloor s/s_1 \rfloor$ geodesics: the same holds for $\Omega_{p_{\lambda C/s}}^C$, since $\Omega_{p_{\lambda C/s}}^{\lambda^2 C} \subset \Omega_{p_{\lambda C/s}}^C$. Now we let $s \to \infty$: then we have $\epsilon \doteq \lambda C/s \to 0$, since $\lambda < \epsilon_s$ but ϵ_s does not increases. It follows that

$$\lim_{s\to\infty} \nu\left(\Omega_{p_\epsilon}^C\right) \ge \lim_{s\to\infty} \lfloor s/s_1 \rfloor = \infty, \tag{3.19}$$

thus the conclusion follows:

$$\lim_{\epsilon\to 0} \nu\left(\Omega_{p_\epsilon}^C\right) = \infty.$$

□

We can summarize as follows: every curve γ passing through 0 such that both $\dot\gamma(0) \notin \Delta_0$ and $\psi(\dot\gamma(0)) + \Delta_0 \notin \Sigma_1 \cup \Sigma_2 \cup \Sigma_3 \cup \Sigma_4$ (for ψ defined in proposition 51) can be realized as an orbit of the one-parameter family of dilations δ_ϵ for a suitable choice of privileged coordinates. Then we prove that along these orbits, the number of bounded geodesics between 0 and $\delta_\epsilon(p)$ (for some p close enough to 0 to be contained in a chart with privileged coordinates) goes to ∞ as $\epsilon \to 0$.

Bibliography

[1] A. A. Agrachev: *Topology of quadratic maps and hessians of smooth maps*, Itogi nauki. VINITI. Algebra. Topologiya. Geometriya, 1988, v.26, 85–124 (Translated in J. Soviet Math. 49, 1990, no. 3, 990Ã©1013).

[2] A. A. Agrachev, U. Boscain, D. Barilari: *Introduction to Riemannian and sub-Riemannian geometry*, lectures notes, available at: https://www.imj-prg.fr/˜davide.barilari/Notes.php

[3] A. A. Agrachev, A. Gentile, A. Lerario: *Geodesics and topology of horizontal-path spaces*, to appear on Geometry & Topology

[4] A. A. Agrachev, J. P. Gauthier: *On the subanalycity of the Carnot-Carathéodory distance*, Ann. I. H. Poincaré – AN 18, 3 (2001) 359–382

[5] A. A. Agrachev, A. Lerario: *Systems of quadratic inequalities*, Proceedings of the London Mathematical Society, (2012) 105 (3).

[6] D. Barilari, U. Boscain, J. P. Gauthier: *On 2-step, corank 2 sub-Riemannian metrics*, SIAM Journal of Control and Optimization, 50 (2012), no. 1, 559-582

[7] A. Bellaïche, *The tangent space in sub-Riemannian geometry*, Journal of Mathematical Sciences, Vol. 83, No. 4. 1997

[8] J. Bochnak, M. Coste, M-F. Roy: *Real Algebraic Geometry*, Springer-Verlag, 1998.

[9] U. Boscain, J. P. Gauthier: *On the Spherical Hausdorff Measure in Step 2 Corank 2 sub-Riemannian Geometry*, arXiv:1210.2615

[10] S. Basu,R. Pollack, M-F. Roy: *Algorithms in Real Algebraic Geometry*, Springer-Verlag, Algorithms and Computation in Mathematics Volume 10

[11] R. Bott: *Nondegenerate critical manifolds*, Annals of Mathematics , Second Series, Vol. 60, No. 2 (Sep., 1954), pp. 248-261

[12] R. Bott: *Lectures on Morse Theory, old and new* Bullettin of the American Mathematical Society, Volume 7, Number 2, September 1982

[13] K. Chang: *Infinite dimensional Morse Theory and Multiple Solution Problems*, Birkhäuser, 1993

[14] Z. Ge: *Horizontal path spaces and Carnot-Carathéodory metrics*, Pacific Journal of Mathematics, volume 161, no.2, 1993

[15] M. Gromov: *Structures Métriques pour les Variétés Riemanniennes*, Cedic-Nathan, Paris (1981)

[16] T. Kato: *Perturbation theory for linear operators*, Springer, 1995

[17] W. Klingenberg: *Lectures on closed geodesics*, Springer, 1978

[18] A. Lerario: *Convex Pencils of real quadratic forms*, Discrete and computational Geometry, Volume 48, Number 4 (2012), 1025-1047.

[19] A. Lerario: *Complexity of intersection of real quadrics and topology of symmetric determinantal varieties*, to appear on Journal of European Mathematical Society

[20] A. Lerario: *On a theorem of Serre in sub-Riemannian geometry*, in preparation

[21] A. Lerario, L. Rizzi: *How many geodesics join two points in a contact sub-Riemannian manifold?*, preprint on arXiv:1405.4294

[22] J. W. Milnor: *Morse theory*, Princeton University Press

[23] R. Montgomery: *A tour of Subriemannian geometries, their geodesics and applications*

[24] L. Nicolaescu: *An invitation to Morse Theory*, Springer Universitext

[25] A. L. Onishchick, E. B. Vinberg: *Lie groups and algebraic groups*, Springer, 1980

[26] R. Palais: *Morse theory on Hilbert manifolds*

[27] P. Pansu *Metriques de Carnot-Carathéodory et quasi isometries des espaces symetriques de rang un*, Annals of Math. 129 , 1-60, 1989

[28] S. Semmes: *An Introduction to Heisenberg Groups in Analysis and Geometry*, Notices of the AMS, volume 50, number 6.

[29] J. P. Serre: *Homologie singulière des espaces fibrés*, Applications, Ann. Math. 54 (1951), 425–505