On the Curvature of Control Systems

Andrei A. Agrachev
SISSA and Steklov Math. Inst.

The curvature tensor on a Riemannian manifold reflects intrinsic properties of the geodesic flow, i.e. the properties which do not depend on the choice of local coordinates. In particular, nonzero curvature is an obstruction for the simultaneous rectification of the geodesics by a change of coordinates.

The knowledge of the curvature gives a fundamental information about the behavior of geodesics and about the structure of optimal synthesis. On the other hand, in order to evaluate the curvature we don't need to solve the geodesics differential equation, it is enough to compute certain polynomials of the partial derivatives of its right-hand side.

Riemannian geodesics are just extremals of a special optimization problem, the problem to find a shortest path between to points on the manifold. Why not to try to find analogues of the curvature for general optimal control problems? This talk is an attempt of a light presentation of the large program going in that direction.

We start from the classical Gaussian curvature of the surface. The curvature at a point of the surface is negative if geodesics started from this point in different directions "diverge". The curvature is positive if these geodesics "converge". How to recognize this divergence – convergence? We draw small geodesic triangles. If the sum of the angles of such a triangle is smaller than π, then the curvature is negative. If the sum of the angles is greater than π, then the curvature is positive. The curvature itself comes from the first meaningful term in the asymptotic expansion of the sum of the angles as the triangles tend to the point.

We can rearrange this construction of the curvature in the geometric control style. Let M be our surface, $TM = \bigcup\limits_{q\in M} T_qM$ be the tangent bundle and $SM = \bigcup\limits_{q\in M} S_qM$ be the "spherical bundle", where $S_qM = \{\xi \in T_qM : |\xi| = 1\}$. A geodesic is uniquely defined by its initial point and velocity. In other words, the geodesics, parametrized by length, define a flow in SM. This geodesic flow is generated by a vector field $\vec{h}$ on SM.

There is another intrinsically defined flow on SM. This flow preserves fibers S_q and is actually the rotation of the circles S_q with a constant velocity. Suppose that a vector field w on SM is the generator of this "vertical" flow. The pair of vector fields $\vec{h}, w$ forms a control system on SM. Geodesic polygons are just projections to M of trajectories of this control system associated to piecewise constant controls. In particular, the expression $e^{t_k\vec{h}} \circ e^{\theta_k w} \cdots \circ e^{t_2\vec{h}} \circ e^{\theta_1 w} \circ e^{t_1\vec{h}}(q)$ presents a geodesic polygon with k edges of the lengths $t_1, \ldots, t_k$ and the angles θ_i between the edges of the numbers i and $i+1$. One can imagine that the curvature can be expressed via the commutator relations of the fields $\vec{h}$ and w. Indeed, the curvatre comes from the first nontrivial commutator relation: it is the coefficient κ in the identity

$$[\vec{h}, [w, \vec{h}]] = \kappa w. \tag{1}$$

It is now clear what to do with more general 2-dimensional problems than the Riemannian ones. Suppose that we have a control system $\dot{q} = f(q, u)$, $u \in U$, and try to minimize a cost $\int\limits_{t_0}^{t_1} \varphi(q(t), u(t))\, dt$. We write the (normal) Hamiltonian of the problem

$$h(\lambda) = \max_{u \in U}(\langle \lambda, f(q, u)\rangle - \varphi(q, u)), \quad \lambda \in T_q^*M,\ q \in M,$$

which is a function on the cotangent bundle. If we are lucky and h is a smooth function, then the Hamiltonian vector field $\vec{h}$ is well-defined and is tangent to the level sets of h. The trajectories of the correspondent Hamiltonian systems are extremals of our optimal control problem. The flow generated by the Hamiltonian system on the cotangent bundle is actually a direct generalization of the Riemannian geodesic flow in spite of the fact that the last flow was originally defined on the tangent bundle. Recall that a Riemannian structure defines a scalar product in the tangent spaces that gives an identification of the tangent and cotangent bundles.

Fix a level set $H = h^{-1}(c)$. Then $H_q = T_q^*M \cap H$ is a curve in the plane T_q^*M. Under certain regularity assumptions such a curve admits a "natural parameter" providing us with a vector field w_q on H_q and, finally, with a vertical vector field w on H. For the vector fields $\vec{h}, w$ we still have the relation (1), where κ is the curvature of the optimal control problem; now, by definition.

This curvature enjoys properties similar to the Gaussian one. The extremals "diverge" and their orbitrarily long segments are locally optimal if κ is negative. They "converge" if κ is positive and a comparison theorem gives estimates for the conjugate points, where extremals cease to be locally optimal.

There are clear restrictions for the given construction of the curvature: it works only for 2-dimensional systems and smooth Hamiltonians. How to manage beyond these restrictions? Let us first discuss the multidimensional case with a smooth Hamiltonian. Here we have the vector field $\vec{h}$ and the fibers H_q. The fiber H_q is a hypersurface in T^*M which, in general, does not admit a natural parametrization. So we have no intrinsically defined vertical vector fields to compute commutators with $\vec{h}$. We have however a vertical vector distribution $D = \{D_\lambda\}_{\lambda \in H}$ of the tangent spaces to the fibers, $D_\lambda = T_\lambda H_q$, $\lambda \in H_q$, $q \in M$. The Hamiltonian flow $e^{t\vec{h}}$ generated by $\vec{h}$ does not preserve the distribution D. The "curvatures" are infinitesimal invariants of the family of distributions $e_*^{t\vec{h}} D$.

There is a nice geometry inside. Take some $\lambda \in H$ and pick 4 subspaces $\left(e_*^{t_1\vec{h}} D\right)_\lambda, \ldots, \left(e_*^{t_4\vec{h}} D\right)_\lambda$ in $T_\lambda H$. Subspaces are points of the correspondent Grassmannian. We may compute the cross-ratio of the given 4 points that is a multidimensional version of the classical cross ratio of 4 points in the projective line. We recover curvatures from the asymptotics of the cross ratio as $t_1, \ldots, t_2$ tend to 0.

Finally, what happens in the case of nonsmooth h? There are basically two very different types of singularities of h which produce singular and bang-bang extremals correspondently. What we are looking for is not a universal set theoretic definition, which can be formally applied everywhere. Rather, we are interested in an object with a clear structure, which could serve to handle specific problems. Because of that we study singular and bang–bang cases separately.

Actually, singular case can be handled similar to the smooth one since singular extremals are normally generated by a smooth easily computable Hamiltonian, though not by h. For the bang–bang case we need a different construction. The work is in progress though some interesting calculations are already done.

One remark to conclude. The curvature is a rather stable invariant and it looks natural to approximate the nonsmooth Hamiltonian by smooth ones, compute the curvature and study its asymptotics as smooth approximations tend to the original nonsmooth object. Actually, this method does not help in the bang–bang case: the cuvature of the approximations usually tends to minus infinity at the switching points. This is not strange: approximating extremals "diverge" very fast, when their velocities tend to the bang–bang values.