Controllability of mechanical control systems

Andrew D. Lewis*

Slide 0

21/04/1995

*Advisor: Richard M. Murray

Notes for Slide 0

Although there has been some work done on control of mechanical systems, a fundamental, systematic study of these systems has not been undertaken. Much of the existing work has relied on the presence of specific structure. The most common examples of the types of structure assumed are symmetry (conservation laws) and constraints. While it may seem counter-intuitive that constraints may help in control theory, this is sometimes in fact the case. The reason is that the constraints provide extra forces (forces of constraint) which can be used to advantage.

In our work, we begin a program to understand the factors which might go into a control theory for general mechanical systems. The most interesting work is done from the Lagrangian perspective where we study systems whose Lagrangians are "kinetic energy minus potential energy." For such systems, the controllability questions are different than those typically asked in nonlinear control theory. In particular, one is often more interested in what happens to *configurations* rather than *states* (which are configurations *and* velocities for these systems). We precisely formulate a new controllability problem, then obtain computable conditions for this new form of controllability in terms of the given structure for the system.

We also have some results for Hamiltonian control systems. These are very clean mathematically, but the restrictions placed on the systems we study in the Hamiltonian setting makes the theory of limited practical value. Nevertheless, it may be possible to combine the Lagrangian and Hamiltonian results to obtain a deeper understanding of each.

Outline

Slide 1

1. A motivating example
2. Background in nonlinear control theory
3. Naive control theory for mechanical systems
4. Lagrangian control systems
5. Examples
6. Hamiltonian control systems

1. A motivating example

The Robotic Leg:

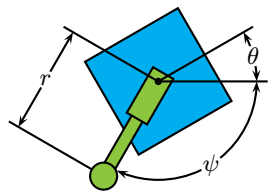

Slide 2

- Inputs are relative torque and leg extension.
- System is "controllable" in the sense that, starting from rest, one can reach any configuration from a given initial configuration.
- As a traditional control system, it is not controllable because of conservation of angular momentum.
- Need a way to account for this "contradiction."

Notes for Slide 2

This simple example exhibits much of the subtle behaviour seen in mechanical control systems. To test one's intuition, it is instructive to look at this example and see if you can predict the structure of the reachable set. Do this with the system starting at rest in some configuration and consider the cases:

1. all inputs are available,
2. only the leg extension is available, and
3. only the relative torque is available.

Is the reachable set of configurations open in Q? Does it contain a neighbourhood of the initial configuration?

2. Background in nonlinear control theory

Consider systems of the form:

$$\dot{x} = f(x) + u^a g_a(x), \qquad x \in M. \tag{NCS}$$

- f is the drift vector field and $g_1, \dots, g_m$ are the control vector fields.

Slide 3

- The reachable set from x_0 in time T is

$$\begin{aligned}\mathcal{R}(x_0,T) = \{x \mid \ & \exists\, c\colon [0,T] \to M \text{ and} \\ & u\colon [0,T] \to \mathbb{R}^m \text{ satisfying (NCS)} \\ & \text{with } c(0) = x_0 \text{ and } c(T) = x\}.\end{aligned}$$

- Also $\mathcal{R}(x_0, \le T) = \bigcup_{0<t\le T} \mathcal{R}(x_0, t)$.

Notes for Slide 3

The drift vector field describes how the system would evolve in the absence of any inputs. Each control vector field specifies a direction in which you can supply actuation.

Observe that since the systems have drift, when we reach the point $c(T)$ we will not remain there if this is not an equilibrium point for f.

2.1. Forms of controllability

Slide 4

- (NCS) is *locally accessible* at x_0 if $\mathcal{R}(x_0, \leq T)$ contains a neighbourhood of M for each $T > 0$.
- (NCS) is *strongly locally accessible* at x_0 if $\mathcal{R}(x_0, T)$ contains a neighbourhood of M for each $T > 0$.
- (NCS) is *small-time locally controllable* (STLC) at x_0 if $\mathcal{R}(x_0, \leq T)$ contains a neighbourhood of x_0.

Notes for Slide 4

We have not been quite precise with our definitions here. The word *local* appear in these definitions for a reason. In the precise definitions for reachable sets, one needs to add the restriction that one not leave a given neighbourhood of the initial point. This has been omitted here for simplicity.

The reason that local accessibility and small-time local controllability are not the same is that the drift may always cause the system to move steadily along in a certain direction, thus never allowing us to get back to the point from which we started. This will be seen in an example on the next slide.

Slide 5

- Example in $\mathbb{R}^2$:

$$\dot{x} = u$$
$$\dot{y} = 1$$

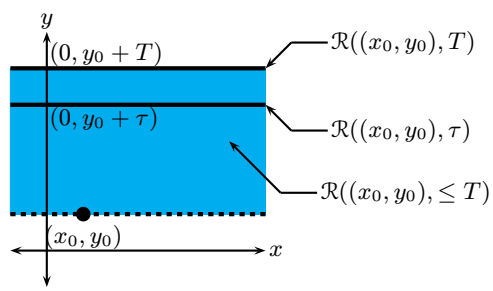

System is locally accessible, but neither strongly locally accessible nor STLC.

Notes for Slide 5

This example, while simple, captures in a fairly precise way the differences between local accessibility, strong local accessibility, and small-time local controllability. To be exact, every nonlinear control system (satisfying regularity conditions) which is locally accessible, but not strongly locally accessible, may be flattened out so that it looks like this example in principle. We also see how the drift term causes the loss of small-time local controllability.

2.2. Distributions and foliations

Slide 6

- A *distribution*, D, on a manifold, M, is a selection of a subspace of possible directions at each point.
- A *foliation*, $\mathscr{F}$, of a manifold, M, is a partitioning of M into disjoint submanifolds. Each disjoint component is called a *leaf*.
- A distribution is *integrable* if the directions selected at each point are tangent to some leaf of a foliation. Denote the foliation by $\mathscr{F}_D$.
- Denote the set of leaves by $M/\mathscr{F}_D$. Call it the *leaf space*.

Slide 7

- Example in $\mathbb{R}^2$: Foliation by circles.

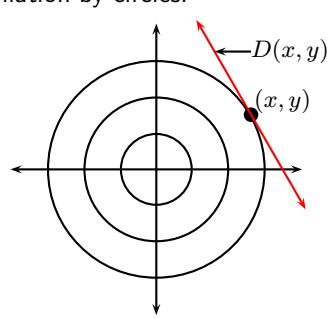

Note that $M/\mathscr{F}_D \simeq \mathbb{R}^+$.

2.3. Lie brackets

Slide 8

- For two vector fields X, Y, their *Lie bracket*, $[X, Y]$, measures the infinitesimal amount that their flows do not commute.
- For the purposes of control theory, Lie brackets of the control and input vector fields generate new "directions" the system may go.
- Sometimes we will speak of a bracket not as a vector field, but as a formal object. This can be done using free Lie algebras.
- If $\mathscr{V}$ is a family of vector fields, $\overline{\mathrm{Lie}}(\mathscr{V})$ will denote the set of all iterated Lie brackets of elements of $\mathscr{V}$. Call this the *involutive closure* of $\mathscr{V}$.
- The distribution defined by $\overline{\mathrm{Lie}}(\mathscr{V})$ is integrable by Frobenius' Theorem.

Notes for Slide 8

A good example of a control system which uses Lie brackets to generate extra motion is a car. Although the only inputs are steering and propulsion, it is nevertheless possible to position the car anywhere in the plane in any orientation. Thus, with two inputs we can generate three directions. The extra direction is a Lie bracket motion.

The *free Lie algebra* in indeterminants $\{X_0, \dots, X_m\}$ may be loosely regarded as the associative polynomial algebra in these indeterminants modulo the relations of satisfying skew-symmetry and Jacobi's identity. The point is that there is a well-defined algebraic structure to the set of brackets so that it makes sense to speak of, for example, $[f, g_a]$ as a bracket rather than as a vector field. This is often convenient and is sometimes necessary to be precise.

2.4. Distributions for control systems

- accessibility distribution is generated by

$$\overline{\mathrm{Lie}}(\{f, g_1, \dots, g_m\}).$$

  Denote it by C.

- strong accessibility distribution is generated by

Slide 9

$$\overline{\mathrm{Lie}}(\{f, g_1, \dots, g_m\}) \setminus \{f\}.$$

  Denote it by C_0.

- Both C and C_0 are integrable.
- Previous example:
  - Accessibility distribution is all of $T\mathbb{R}^2$.
  - Strong accessibility distribution is generated by $\frac{\partial}{\partial x}$.

Notes for Slide 9

The leaves of the foliation defined by the accessibility distribution are, by construction, invariant under the control system (NCS). The leaves defined by the strong accessibility distribution may not be invariant in the sense that the trajectories of the control system will remain on the leaf on which they started. This is a consequence of the drift vector field not necessarily being contained in the strong accessibility distribution.

However, the strong accessibility distribution does satisfy a weaker, infinitesimal, version of invariance. This is manifested in the following fact: Let x_1, x_2 be two points on the same leaf of the foliation defined by the strong accessibility distribution. If the control system evolves under any set of inputs for some time T, then x_1, x_2 will flow to points on the same leaf of the strong accessibility distribution. Thus the leaves of the foliation defined by the strong accessibility distribution are not necessarily invariant, but they flow along under the control system from one leaf to another.

This behaviour is manifested in the simple example presented in Section 2.1. There we can see that horizontal lines, which are the leaves of the strong accessibility distribution, flow to horizontal lines. Note that within each leaf, the system is strongly locally accessible.

The behaviour of leaves moving from one to another gives rise to the notion of dynamics on the leaf space. Dynamics on the leaf space will come back to us later in the talk.

2.5. Controllability tests

- If $\dim(C(x_0)) = \dim(M)$ then (NCS) is locally accessible at x_0.
- If $\dim(C_0(x_0)) = \dim(M)$ then (NCS) is strongly locally accessible at x_0.
- The finest known sufficient condition for STLC[1] is hard to say. We state a simplified version:

Slide 10

  - Call a bracket "bad" if it has an odd number of f's in it and an even number of each of the g_a's.
  - Otherwise call it "good."
  - If every bad bracket can be written as a linear combination of good brackets of lower order, and if the system is locally accessible at x_0, then (NCS) is STLC at x_0.

[1] Sussmann, H. J., *A general theorem on local controllability*, SIAM Journal on Control and Optimization, **25**, 158–194, 1987.

Notes for Slide 10

Under regularity conditions, the rank conditions on C and C_0 are necessary as well as sufficient for local accessibility and strong local accessibility, respectively.

Useful necessary and sufficient conditions for STLC are not known.

2.6. Decompositions

- Use Frobenius' Theorem to obtain decompositions when (NCS) is not locally accessible or strongly locally accessible.
- For accessibility distribution:

Slide 11

$$\begin{aligned}\dot{x}_1 &= f_1(x_1, x_2) + u^a g_a(x_1, x_2)\\ \dot{x}_2 &= 0\end{aligned}$$

  For fixed x_2, the top system is locally accessible.

- For strong accessibility distribution:

$$\begin{aligned}\dot{\bar{x}}_1 &= \bar{f}_1(\bar{x}_1, \bar{x}_2) + u^a g_a(\bar{x}_1, \bar{x}_2)\\ \dot{\bar{x}}_2 &= \bar{f}_2(\bar{x}_2)\end{aligned}$$

  For fixed $\bar{x}_2$, the top system is strongly locally accessible.

Notes for Slide 11

Recall that Frobenius' Theorem tells us that if D is a distribution of constant rank k on M, then there are coordinates $(x^1, \dots, x^k, x^{k+1}, \dots, x^n)$ for M so that

$$D(x) = \langle \frac{\partial}{\partial x^1}, \dots, \frac{\partial}{\partial x^k} \rangle_{\mathbb{R}}.$$

Note that fixing x_2 in the local accessibility decomposition is meaningful. It means fixing a leaf of the foliation defined by the accessibility distribution. However, fixing $\bar{x}_2$ for the strong local accessibility decomposition only has meaning when the $\bar{x}_2$ component of the drift vector field is zero. This occurs for exactly those problems where the drift vector field is in C_0.

3. Naive control theory for mechanical systems

- When the velocities of the system are assumed controlled:

$$\dot{q} = u^a X_a(q),$$

we have the *kinematic* or *nonholonomic* problem.

Slide 12

- It is easy to describe the reachable set for this problem.
- Add dynamical effects in the following simple manner:

$$\begin{aligned}\dot{q} &= v\\ \dot{v} &= u^a X_a(q).\end{aligned}$$

- This integrates the inputs before they reach the configurations and may be regarded as an *extremely* crude model of dynamical effects.

Notes for Slide 12

To describe the reachable sets for the nonholonomic problem, simply compute the leaves of the foliation corresponding to the integrable distribution generated by $\overline{\mathrm{Lie}}(\{X_1, \dots, X_m\})$.

The technique of integrating the inputs once does not change the controllability as far as the configurations are concerned. What was reachable before is still reachable.

The main problem with integration of the inputs in the given naive fashion is contained in the fact that this is not a coordinate invariant process. To make the process coordinate invariant, one needs to add further structure. As we shall see, this structure is kinetic energy by way of a Riemannian metric. This provides an important mechanism for the breakdown of the naive setup presented here: Coriolis forces. We shall see that there are other factors which can also cause the naive approach to fail. The bottom line is that the naive approach is far too naive to be close to accurate in anything like a general setting. There is simply too big a difference between supplying torques and specifying velocities.

One may want to think about our robotic leg example in the context of what we have said here. Do you think that the above model of system dynamics will capture the correct behaviour of the leg?

I have also seen the controls integrated, i.e.,

$$\begin{aligned}\dot{q} &= v^a X_a(q)\\ \dot{v}^a &= u^a.\end{aligned}$$

However, I don't understand this approach.

4. Lagrangian control systems

- Consider Lagrangian to be "kinetic energy minus potential energy."
- This means a Riemannian metric, g, (given by the inertia matrix) on the configuration space, Q, and a potential function on Q.
- Input forces are modelled as one-forms, $F^1, \dots, F^m$, on Q.

Slide 13

- We will be interested in the vector fields, $Y_1, \dots, Y_m$, on Q obtained by "multiplying the F's by the inverse of the inertia matrix."
- The control system is then

$$\ddot{q}^i + \Gamma^i_{jk}\dot{q}^j\dot{q}^k + g^{ij}\frac{\partial V}{\partial q^j} = u^a Y^i_a. \qquad \text{(LCS)}$$

The Γ^i_{jk}'s are the *Christoffel symbols* and are computable from the inertia matrix.

Notes for Slide 13

We shall deem this to be a "natural" class of Lagrangian systems. By natural we mean that we want to be able to describe the system in terms of the given structure, the Riemannian metric and the potential function in this case. And, of course, the class of systems whose Lagrangians are kinetic energy minus potential energy forms a large and physically interesting group.

Note that the input vector fields, $Y_1, \dots, Y_m$, *do* depend on the metric in the sense that they are obtained from the one-forms $F^1, \dots, F^m$ by "sharping" them under the metric. It is the one-forms, however, which describe the forces in a fundamental way. The input vector fields are what appear in our computations. Also notice that the gradient of the potential energy enters into (LCS) and it is this gradient vector field which will represent the effects of the potential energy in our calculations.

The assumption that the forces be represented as stated really means that we are restricting the directions in which we can apply forces to be functions of configuration. In general, one may wish to allow these directions to be functions of time and velocity.

Slide 14

4.1. Controllability definitions for Lagrangian control systems

- The *reachable set* from q_0 in time T is

$$\begin{aligned}\mathcal{R}_Q(q_0,T) = \{q \in Q \mid \ & \exists\, c\colon [0,T] \to Q \text{ and} \\ & u\colon [0,T] \to \mathbb{R}^m \text{ satisfying (LCS) with} \\ & c'(0) = (q_0, 0) \text{ and } c'(T) \in T_qQ\}.\end{aligned}$$

- Also

$$\mathcal{R}_Q(q_0, \leq T) = \bigcup_{0<t\leq T} \mathcal{R}_Q(q_0, t).$$

Notes for Slide 14

Notice that we ask the system to start at zero velocity, but when it reaches its final configuration it may be moving with some arbitrary, non-zero velocity.

If we were to cast (LCS) as a standard nonlinear control system, the usual definitions of the reachable set look like "What is the set of configurations *and velocities* that can be reached from a given initial configuration *and velocity*?" We have done something different, and perhaps more natural.

Slide 15

- (LCS) is *locally configuration accessible* at q_0 if $\mathcal{R}_Q(q_0, \leq T)$ contains a neighbourhood of Q for each $T > 0$.
- (LCS) is *strongly locally configuration accessible* at q_0 if $\mathcal{R}_Q(q_0, T)$ contains a neighbourhood of Q for each $T > 0$.
- (LCS) is *small-time locally configuration controllable* (STLCC) at q_0 if $\mathcal{R}_Q(q_0, \leq T)$ contains a neighbourhood of q_0.

Notes for Slide 15

Note that definitions are made in terms of *configurations*. If we cast (LCS) as a standard nonlinear control system, the usual definitions are made in terms of configurations *and* velocities. It is the configurations in which one is typically interested. It should be emphasised at this point that these definitions are *new*. This may be surprising as they reflect an obvious, mechanically interesting question.

4.2. Symmetric products

- *Covariant derivative* of two vector fields:

$$(\nabla_X Y)^i = \frac{\partial Y^i}{\partial q^j} X^j + \Gamma^i_{jk} X^j Y^k.$$

- Define a *symmetric product* on the set of vector fields by

Slide 16

$$\langle X : Y \rangle = \nabla_X Y + \nabla_Y X.$$

This depends on the Riemannian metric.

- $\langle X : Y \rangle$ is a vector field.
- If $\mathscr{V}$ is a set of vector fields on Q, define $\overline{\mathrm{Sym}}(\mathscr{V})$ to be the set of all iterated symmetric products of elements of $\mathscr{V}$. Call this the *symmetric closure* of $\mathscr{V}$.

Notes for Slide 16

The covariant derivative takes two vector fields and returns a vector field. Thus the symmetric product does the same. It is also clearly symmetric. It is worth noting that this product depends upon the Riemannian structure, as the Lie bracket relies only upon the differentiable structure of the manifold.

Just as it is possible to speak of free Lie algebras, it is possible to speak of "free symmetric algebras." Such objects allow one to speak of a symmetric product as an object rather than as a vector field and we shall sometimes do this.

One should note the similarities and differences between the symmetric closure and the involutive closure. The biggest difference is that the involutive closure defines an integrable distribution, but the same cannot be guaranteed for the symmetric closure.

4.3. Distributions of symmetric products and Lie brackets

- We wish to compute the accessibility distribution, C, at points of zero velocity, denoted $(q, 0)$.
- At $(q, 0)$ define *horizontal* to be in the "q-direction" and *vertical* to be in the "v-direction."

Slide 17

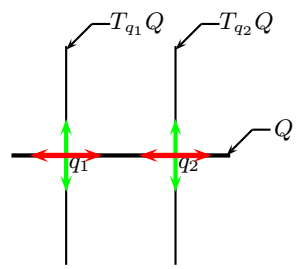

Notes for Slide 17

The reason that we are interested in the accessibility distribution at points of zero velocity is that these are the initial conditions in our versions of controllability. Thus the distribution evaluated at these points will be tangent to the reachable set.

We shall describe the accessibility distribution in terms of two distributions on Q. One distribution will represent the vertical directions in the figure, and the other the horizontal directions.

Please note that the decomposition into horizontal and vertical as described above does not, in general, make sense except at points where the velocity is zero. Other decompositions may be made at points of non-zero velocity, but we shall not go into that here.

- Denote $\mathscr{Y} = \{Y_1, \dots, Y_m\}$.
- All vertical directions for C are generated by $\mathbb{R}$-linear combinations of vector fields from $\overline{\mathrm{Sym}}(\mathscr{Y} \cup \{\mathrm{grad}\, V\})$.
- All horizontal directions for C are generated by $\mathbb{R}$-linear combinations of vector fields from $\overline{\mathrm{Lie}}(\overline{\mathrm{Sym}}(\mathscr{Y} \cup \{\mathrm{grad}\, V\}))$.

Slide 18

- An algorithm determines which $\mathbb{R}$-linear combinations to choose.
- Define two distributions on Q,

$$C_{\mathit{ver}}(\mathscr{Y}, V),$$
$$C_{\mathit{hor}}(\mathscr{Y}, V)$$

as the vertical and horizontal directions, respectively, for C.

Notes for Slide 18

Although we have an algorithm for computing which vector fields from $\overline{\mathrm{Sym}}(\mathscr{Y} \cup \{\mathrm{grad}\, V\})$ and $\overline{\mathrm{Lie}}(\overline{\mathrm{Sym}}(\mathscr{Y} \cup \{\mathrm{grad}\, V\}))$ are present in the accessibility distribution, it would be nice to compute an inductive formula for these vector fields. This seems to be quite difficult.

Slide 19

- It is easy to show that $C_{hor}(\mathscr{Y}, V)$ is an integrable distribution.
- Special Case of $V = 0$:
  - The vertical directions are generated by $\overline{\mathrm{Sym}}(\mathscr{Y})$.
  - The horizontal directions are generated by $\overline{\mathrm{Lie}}(\overline{\mathrm{Sym}}(\mathscr{Y}))$.
  - $C_{ver}(\mathscr{Y}, V) \subset C_{hor}(\mathscr{Y}, V)$.

Notes for Slide 19

The integrability of $C_{hor}(\mathscr{Y}, V)$ is important in our decompositions we state later.

The fact that $C_{ver}(\mathscr{Y}, V) \subset C_{hor}(\mathscr{Y}, V)$ means that solutions starting with zero velocity will remain on leaves of the foliation defined by $C_{hor}(\mathscr{Y}, V)$. Note that this only happens in general if $V = 0$. There are other special cases when the solutions will remain on the leaves of $C_{hor}(\mathscr{Y}, V)$. This is discussed further when we present decompositions for Lagrangian control systems.

Slide 20

4.4. Controllability tests for mechanical control systems

- If $\dim(C_{hor}(\mathscr{Y}, V)(q_0)) = \dim(Q)$ then (LCS) is locally configuration accessible at q_0.
- Conditions for STLCC are similar to those for STLC in (NCS).
  - We shall call a symmetric product in $\overline{\mathrm{Sym}}(\mathscr{Y} \cup \{\mathrm{grad}\, V\})$ "bad" if the number of times that each vector field Y_a appears is even.
  - Otherwise the symmetric product is called "good."
  - If each bad symmetric product can be written as a linear combination of good symmetric products of lower order, and if the system is locally configuration accessible at q_0, then (LCS) is STLCC at q_0.

Notes for Slide 20

Notice that we do not bother with strong local configuration accessibility. This is something that we would like to understand better.

The conditions for STLCC actually give more than is implied. It is, in fact, possible to go from a configuration q_0 at rest to a neighbourhood of q_0 *at rest*. This allows us to define the notion of "equilibrium controllability" whereby one can steer between two equilibrium points at rest.

4.5. Decompositions for Lagrangian systems

- In the case when a Lagrangian system is *not* locally configuration accessible, what does it look like?
- Use Frobenius' Theorem for the integrable distribution $C_{hor}(\mathscr{Y}, V)$.
- We find coordinates $q = (x, y)$ for Q so that

Slide 21

$$\ddot{x}^i + \Gamma^i_{jk}(x,y)\dot{x}^j\dot{x}^k + \Gamma^i_{j\alpha}(x,y)\dot{x}^j\dot{y}^\alpha + \Gamma^i_{\alpha\beta}(x,y)\dot{y}^\alpha\dot{y}^\beta + g^{ij}\frac{\partial V}{\partial x^j} + g^{i\alpha}\frac{\partial V}{\partial y^\alpha} = u^a Y^i_a$$

$$\ddot{y}^\alpha + \Gamma^\alpha_{j\beta}(x,y)\dot{x}^j\dot{y}^\beta + \Gamma^\alpha_{\beta\gamma}(x,y)\dot{y}^\beta\dot{y}^\gamma + g^{\alpha j}\frac{\partial V}{\partial x^j} + g^{\alpha\beta}\frac{\partial V}{\partial y^\beta} = 0$$

- For fixed y and with $\dot{y} = 0$, the top system is locally configuration accessible.

Notes for Slide 21

The essential points of interest in the decomposition are that there are no inputs present in the bottom system, and no term $\Gamma^\alpha_{ij}\dot{x}^i\dot{x}^j$ present in the bottom system. Note that if the gradient of the potential is in the involutive closure of the inputs, then no potential terms will appear in the bottom system.

The fixing of y and setting $\dot{y} = 0$ is meaningful when $V = 0$. It means consider the Lagrangian control system determined by pulling back the Riemannian metric to a leaf of $C_{hor}(\mathscr{Y}, V)$. This process is well-defined and, in each case, produces a locally configuration accessible mechanical system.

We see here how the system may not remain on leaves of $C_{hor}(\mathscr{Y}, V)$ even if it starts with zero initial velocity. The potential terms in the bottom equation will take effect and move us off the leaf on which we started.

5. Examples

Robotic leg Inputs are Y_1 (relative torque) and Y_2 (leg extension).

Case 1. Y_1 and Y_2:

- $C_{hor}(\mathscr{Y}, V) = \langle Y_1, Y_2, [Y_1, Y_2] \rangle = TQ$.
- $2\nabla_{Y_1} Y_1 (= \nabla_{Y_1} Y_1 + \nabla_{Y_1} Y_1)$ is a multiple of Y_2.
- System is STLCC.

Slide 22

Case 2. Y_1:

- $C_{hor}(\mathscr{Y}, V) = \langle Y_1, 2\nabla_{Y_1} Y_1, [Y_1, 2\nabla_{Y_1} Y_1] \rangle = TQ$.
- $2\nabla_{Y_1} Y_1$ does not lie in the span of the inputs.
- System is locally configuration accessible.
- System may not be (in fact isn't) STLCC.

Notes for Slide 22

The symmetric product $2\nabla_{Y_1} Y_1$ is bad since Y_1 appears in it twice. This is the only type of bad symmetric product we shall encounter in this talk.

Observe that the "unSTLCCness" in Case 2 is a consequence of the presence of "Coriolis forces" which push the ball increasingly outward no matter what else is happening to the controlled variables.

We should be clear, however, in saying that only the sufficient conditions for STLCC have been violated. Only because we know more about the dynamics of this simple system are we able to say that it is, in fact, not STLCC.

Case 3. Y_2:

- $C_{hor}(\mathscr{Y}, V) = \langle Y_2 \rangle \subsetneq TQ$.
- System is not locally configuration accessible.
- The local configuration accessibility decomposition of the system is given by

Slide 23

$$\ddot{r} - r\dot{\psi}^2 = \frac{1}{m}u_1$$

$$\ddot{\theta} = 0$$

$$\ddot{\psi} + \frac{2}{r}\dot{r}\dot{\psi} = 0.$$

Notes for Slide 23

The readers should assure themselves that the decomposition has the advertised properties: No inputs appear in the bottom system and there are no terms quadratic in $\dot{r}$ in the bottom system. In this case, since the potential energy is zero, this means that θ and ψ remain at their initial values if $\dot{\theta}(0)$ and $\dot{\psi}(0)$ are both zero.

Planar rigid body The system is

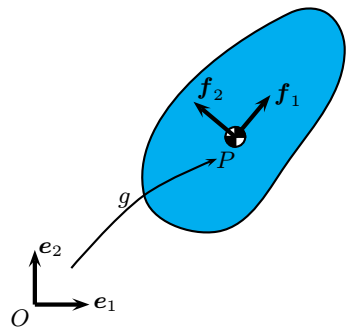

Slide 24

The configuration space is $Q = SE(2)$ with coordinates (x, y, θ). The inputs are

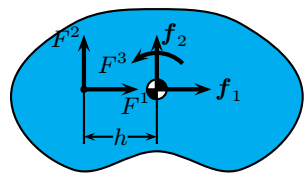

Notes for Slide 24

The coordinates (x, y) denote the position of the centre of mass, P, with respect to the inertial reference frame with origin at O, and the angle θ denotes the orientation of the body frame with respect to the inertial frame.

This system is an example of a class of systems on Lie groups where the Lagrangian is left-invariant (here defined by a left-invariant Riemannian metric), and the inputs are also left-invariant. The inputs are then simply described as a set of elements of the dual of the Lie algebra which are then pushed-forward under left translation to give a left-invariant one-form.

Slide 25

Input vector fields are Y_1 (force directed towards the centre of mass), Y_2 (force directed perpendicular to the direction of the centre of mass), and Y_3 (torque at the centre of mass).

Case 1. Y_1 and Y_2:

- $C_{hor}(\mathscr{Y}, V) = \langle Y_1, Y_2, [Y_1, Y_2] \rangle = TQ$.
- $2\nabla_{Y_2} Y_2$ is in the span of the inputs.
- System is STLCC.

Case 2. Y_1 and Y_3:

- $C_{hor}(\mathscr{Y}, V) = \langle Y_1, Y_3, [Y_1, Y_3] \rangle = TQ$.
- All covariant derivative terms are zero.
- System is STLCC.

Slide 26

Case 3. Y_2 and Y_3:

- $C_{hor}(\mathscr{Y}, V) = \langle Y_2, Y_3, [Y_2, Y_3] \rangle = TQ$.
- $2\nabla_{Y_2} Y_2$ is not in the span of the inputs.
- System is locally configuration accessible.
- System may not be (but actually is) STLCC.

Case 4. Y_2:

- $C_{hor}(\mathscr{Y}, V) = \langle Y_2, \nabla_{Y_2} Y_2, [Y_2, \nabla_{Y_2} Y_2] \rangle = TQ$.
- $2\nabla_{Y_2} Y_2$ is not in the span of the inputs.
- System is locally configuration accessible.
- System may not be STLCC.

Notes for Slide 26

That the system is STLCC in Case 3 is a consequence of the fact that the force-couple pair may be transfered so that the force may be equivalently applied at the centre of mass. This equivalent system is trivially STLCC.

Notice that in Case 4 the "unSTLCCness" cannot come from Coriolis terms (the metric is flat), but rather must come from non-trivial interaction of the inputs. Thus we have seen two ways in which our naive setup presented in Section 3 may fail for mechanical systems with accurate dynamical effects.

In this example, unlike the example of the robotic leg, we are not sure that the system is not STLCC. We believe, however, that it is indeed not STLCC.

Slide 27

Case 4. Y_1:

- $C_{hor}(\mathscr{Y}, V) = \langle Y_1 \rangle \subsetneq TQ$.
- System is not locally configuration accessible.
- The local configuration accessibility decomposition is rendered in coordinates

$$\xi = x\cos\theta + y\sin\theta,$$
$$\eta = -x\sin\theta + y\cos\theta, \quad \psi = \theta.$$

We have

$$\ddot{\xi} + 2\left(\frac{m\eta^2}{J} - \frac{J + m\eta^2}{J}\right)\dot{\eta}\dot{\psi} +$$
$$\left(\frac{m\xi\eta^2}{J} - \frac{\xi J + m\xi\eta^2}{J}\right)\dot{\psi}^2 = \left(\frac{J + m\eta^2}{J} - \frac{\eta^2}{J}\right)u_1$$

Slide 28

$$\ddot{\eta} + 2\left(\frac{J + m\xi^2}{J} - \frac{m\xi^2}{J}\right)\dot{\xi}\dot{\psi} +$$
$$\left(\frac{m\eta\xi^2}{J} - \frac{\eta J + m\eta\xi^2}{J}\right)\dot{\psi}^2 = 0$$
$$\ddot{\psi} = 0$$

Notes for Slide 28

Observe that the coordinates for the local configuration accessibility decomposition are rotated to be aligned with the body frame for the system. It makes sense that these be the coordinates we choose since the direction Y_1 will not "change" with θ in these coordinates. Also notice that the equations are a lot messier in these coordinates. In spite of this, it is easier to see what is going on from the point of view of configuration controllability.

It is interesting to note that the author found the coordinates not by intuition, but by solving the simple PDE's that one gets by trying to find Frobenius coordinates. The intuition came later.

It is also worth noting that neither the robotic leg nor the planar rigid body is linearly controllable with any of the given inputs.

Pendulum on a cart The system is

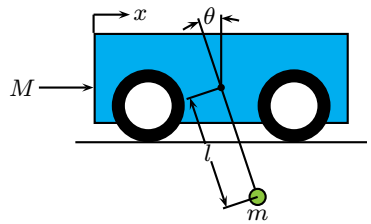

Slide 29 Use coordinates (x, θ). The input is denoted Y_1 and is the force pushing the cart along the surface. The potential function is specified by gravity.

Case 1. $\theta \notin \{0, \pi\}$:

- $C_{hor}(\mathscr{Y}, V) = \langle Y_1, \nabla_{Y_1} Y_1 \rangle = TQ$.
- $2\nabla_{Y_1} Y_1$ is not a multiple of the input.
- System is locally configuration accessible.
- System may not be STLCC.

Slide 30

Case 2. $\theta \in \{0, \pi\}$:

- $C_{hor}(\mathscr{Y}, V) = \langle Y_1, \nabla_{\operatorname{grad} V} Y_1 + \nabla_{Y_1} \operatorname{grad} V \rangle = TQ$.
- $2\nabla_{Y_1} Y_1 = 0$.
- System is STLCC.

Notes for Slide 30

At the points where $\theta \in \{0, \pi\}$, the linearisation is controllable so it must be that the nonlinear system is STLC, and hence STLCC.

5.1. Wrap up for Lagrangian control theory

Slide 31

- Conclusions:
  - *New and meaningful* definitions of controllability for simple mechanical control systems.
  - Computable conditions for these new notions of controllability in terms of data defined on Q.
  - Surprising geometric insight.
- Future explorations:
  - Why covariant derivatives?
  - Find a smaller generating set.
  - What about non-zero initial velocities?
  - More general Lagrangians.
  - Better conditions for STLCC.
  - Synthesis.

Notes for Slide 31

It cannot be emphasised enough that the definitions for controllability we present, while natural enough, have not appeared in the literature. This is especially interesting as we are able to derive fairly nice conditions for checking these definitions of controllability.

We say that the geometric insights gained by this work are surprising since the appearance of the covariant derivative is not something which was expected at the outset. Although we have seen a little of what the covariant derivative terms give us with the centrifugal effects in the robotic leg, we still do not fully understand why they are present.

Notice that in the definitions of $\mathscr{C}_{\mathsf{ver}}(\mathscr{Y}, V)$ and $\mathscr{C}_{\mathsf{hor}}(\mathscr{Y}, V)$, we added *all* covariant derivatives and Lie brackets. In fact we do not need all of them since they will not be linearly independent in general. It would be helpful to determine a way of computing a generating set which requires the computation of fewer covariant derivatives and Lie brackets. This is related to finding Philip Hall bases for free Lie algebras.

6. Hamiltonian control systems

Slide 32

- A Hamiltonian system is specified by a Hamiltonian function (plus things described later) which is the energy of the system.
- A Hamiltonian *control* system is specified by the *system Hamiltonian* H_0 plus *control Hamiltonians*, $H_1, \dots, H_m$.
- Associated with a Hamiltonian function H is the Hamiltonian vector field X_H. A *Hamiltonian control system* has the form

$$\dot{x} = X_{H_0} + u^a X_{H_a}. \qquad \text{(HCS)}$$

Notes for Slide 32

As with Lagrangian systems, we choose a "natural" structure for Hamiltonian control systems. While it is fairly clear that the Lagrangian setup represents a lot of interesting systems, the same cannot so easily be said of the Hamiltonian control framework we present. It *is*, however, true that a lot of systems fall into the given framework, at least locally. The reason for this is that, for many systems, the control Hamiltonians are simply coordinate functions in some chart and so the corresponding vector field is, by construction, locally Hamiltonian. The global issues of this vector field not being properly Hamiltonian are not so important to us since we are in the business of presenting local decomposition theorems.

6.1. Geometry of Hamiltonian systems (Poisson manifolds)

- The general setting for Hamiltonian mechanics is a Poisson manifold.
- A *Poisson manifold* may (loosely) be characterised as a manifold having local coordinates $(q^1, \dots, q^n, p_1, \dots, p_n, c^1, \dots, c^s)$ so that the Hamiltonian vector field with Hamiltonian H is given by

Slide 33

$$\begin{aligned} \dot{q}^i &= \frac{\partial H}{\partial p_i} \\ \dot{p}_i &= -\frac{\partial H}{\partial q^i} \\ \dot{c}^a &= 0 \end{aligned}$$

- The surfaces defined by $c^a = const.$ are invariant under all Hamiltonian vector fields.
- A Poisson manifold P is *symplectic* if $\dim(P) = 2n$. Thus a symplectic manifold is a "nondegenerate" Poisson manifold.

Notes for Slide 33

Our "definition" of a Poisson manifold is backwards. The real definition defines a tensor field on P of a certain type. With this structure, one may move on to define Hamiltonian vector fields. However, the picture of a Poisson manifold being the disjoint union of a bunch of symplectic manifolds is a general one.

When we define a Hamiltonian control system we shall assume that it evolves on a *symplectic* manifold. Poisson manifolds come up later in the analysis.

Slide 34

6.2. Decompositions for Hamiltonian control systems

- Let C be the accessibility distribution and let C_0 be the strong accessibility distribution. Denote the corresponding foliations by $\mathscr{F}_C$ and $\mathscr{F}_{C_0}$.
- Study dynamics on the leaves and on the leaf space.
- On each leaf of $\mathscr{F}_C$ the symplectic structure is degenerate. "Quotienting out" the degeneracy gives a "reduced" symplectic manifold.
- The system (HCS) drops to this reduced manifold and there gives a locally accessible Hamiltonian control system.
- There is one such reduced control system for each leaf of $\mathscr{F}_C$.
- The reduced dynamics may be regarded as the "locally accessible dynamics."

Notes for Slide 34

The distributions C and C_0 are computed in the usual manner described above for general nonlinear systems. Because the drift and control vector fields are Hamiltonian, there is a lot of structure to their involutive closure.

This "quotienting out" procedure is very much the same as the reduction of the level sets of the momentum map in classical Hamiltonian reduction. In fact, in many examples it is *exactly* this procedure: the control system leaves some conservation law undisturbed and so restricts to the level sets of the conserved quantities. The reduction then follows the classical methods.

Observe that we get many sets of locally accessible dynamics, one for each leaf of $\mathscr{F}_C$.

- For the strong accessibility distribution we directly consider the quotient $P/\mathscr{F}_{C_0}$. This turns out to have a natural Poisson structure.
- Since the leaves of $\mathscr{F}_{C_0}$ are invariant under the control vector fields, these vector fields vanish on the quotient.
- Since the drift vector field may not leave these leaves invariant, it may drop non-trivially to the quotient.

Slide 35

- The vector field so obtained describes the strongly locally inaccessible dynamics.
-

$$\begin{tikzcd} & P \supset \Lambda \arrow[dl] \arrow[dr] & \\ P/\mathscr{F}_{C_0} & & \Lambda/\ker(\Omega) \end{tikzcd}$$

Notes for Slide 35

Note the contrast with the accessibility decomposition. There is only one representation of the strongly locally inaccessible dynamics whereas, for the locally accessible dynamics, there is one representation for each leaf of $\mathscr{F}_C$.

Note also that in the problems where the drift vector field is contained in C_0, the strongly locally inaccessible dynamics will be trivial.

To summarise the Hamiltonian decompositions, there is a family of locally accessible Hamiltonian control systems, one for each leaf $\mathscr{F}_C$. The dynamics of this locally accessible system represent the accessible Hamiltonian dynamics on that leaf. For the whole system there is a single (generalised) Hamiltonian system which represents the strongly locally inaccessible dynamics. The fact that these reduced systems are (generalised) Hamiltonian is a consequence of our choosing the inputs to be Hamiltonian vector fields.

The examples as Hamiltonian control systems

Robotic leg with both inputs

- Hamiltonian is

Slide 36

$$H = \frac{1}{2J}p_\theta^2 + \frac{1}{2m}\left(p_r^2 + r^{-2}p_\psi^2\right).$$

- Control Hamiltonians are $H_1 = \theta - \psi$ and $H_2 = r$.
- Leaves of C are given by $\text{ang. mom.} = p_\theta + p_\psi = const. = \mu$.
- Reduced manifold is described by coordinates $(r, \phi = \theta - \psi, p_r, p_\phi = p_\theta - p_\psi)$.

Notes for Slide 36

The calculations for Hamiltonian control systems need much more information than we have presented here. Therefore, we shall present just the "answers" and not go into how to compute them in any detail.

Note that $\theta - \psi$ is not really a function on Q and so X_{H_1} is only locally Hamiltonian.

The reduced control system lives on $T^*(\mathbb{S}^1 \times \mathbb{R})$ with its canonical symplectic structure. One can obtain this reduction by classical methods of group reduction.

- Reduced Hamiltonian is

$$H_\mu = \frac{1}{2m}p_r^2 + \left(\frac{1}{8J} + \frac{1}{8mr^2}\right)p_\phi^2 + \left(\frac{1}{4mr^2} - \frac{1}{4J}\right)\mu p_\phi + \left(\frac{1}{8J} + \frac{1}{8mr^2}\right)\mu^2.$$

Slide 37

- Reduced control Hamiltonians are $H_{1,\mu} = \phi$ and $H_{2,\mu} = 0$.
- Strongly inaccessible dynamics are trivial.

Robotic leg with torque between bodies

- Same as both inputs.

Notes for Slide 37

It may be verified that the Hamiltonian control system on $T^*(\mathbb{S}^1 \times \mathbb{R})$ with Hamiltonian H_μ and control Hamiltonian $H_{1,\mu}$ is indeed locally accessible.

Robotic leg with leg extension

- Leaves of C are given by $p_\theta = const. = \mu$ and $p_\psi = const. = \nu$.
- Reduced manifold is described by coordinates (r, p_r).
- Reduced Hamiltonian is

Slide 38

$$H_{\mu,\nu} = \frac{1}{2m} p_r^2 + \frac{1}{2mr^2}\nu^2 + \frac{1}{2J}\mu^2.$$

- Reduced control Hamiltonian is $H_{2,\mu,\nu} = r$.
- The strongly inaccessible dynamics live on a Poisson manifold with coordinates $(\theta, p_\theta, p_\psi)$.
- The Hamiltonian is $\tilde{H} = \frac{1}{2J} p_\theta^2$.

Notes for Slide 38

In this case, the angular momentum of the body and of the leg are both conserved by the input. For the other inputs, only their sum is conserved. The reduction in the case of the leg extension input can also be accomplished using group methods, this time with the group $\mathbb{T}^2$.

The coordinates $(\theta, p_\theta, p_\psi)$ are coordinates for a Poisson manifold where $q_1 = \theta$, $p_1 = p_\theta$, and $c_1 = p_\psi$. The strongly locally inaccessible dynamics represent the fact that, with this input, the rigid body part of the system moves unaffected by the inputs.

Planar rigid body

- Cannot be represented as a Hamiltonian control system.

Slide 39

Pendulum on a cart

- Can be represented as a Hamiltonian control system but it is trivial since the system is strongly locally accessible.

Notes for Slide 39

The planar rigid body cannot be represented as a Hamiltonian control system since the input vector fields Y_1 and Y_2 are not even locally Hamiltonian. The vector field Y_3 is locally Hamiltonian so the system can be locally represented as a Hamiltonian system with only this input. But this is the most uninteresting input.

For the pendulum on the cart, the system is strongly locally accessible so the locally accessible dynamics are the original dynamics and the strongly locally inaccessible dynamics are trivial.

Slide 40

6.3. Wrap up for Hamiltonian control theory

- Conclusions:
  - The "controllable" dynamics are Hamiltonian.
  - The "uncontrollable" dynamics are Hamiltonian.
- Future explorations:
  - Make connections with Lagrangian results.
  - More general inputs.

Notes for Slide 40

Although this Hamiltonian picture is very attractive, it may not offer so much from a practical viewpoint. The first restriction is that it needs input vector fields which are Hamiltonian. This is a large restriction. If this could be fixed up somehow it would be very interesting. Also, for Hamiltonian control systems on cotangent bundles, this presentation does not address the important mechanical issues, some of which are touched upon in the Lagrangian control theory we presented.