Notes on non-holonomic constraints

by Flip Tanedo

For P3318, Spring 2013

Provided **for culture**, new information is non-examinable. Corrections? pt267@cornell.edu.

Most of the discussion here draws from the references at the end of the document.

1 Types of constraints

1. **Holonomic**: $f(q_1, \cdots, q_n, t) = 0$. Such a constraint can be used to reduce the number of degrees of freedom in a system.
2. **Non-holonomic**: $f(q_1, \cdots, q_n, \dot{q}_1, \cdots, \dot{q_n}, t) = 0$.
3. **Neither**: not described by equations, for example $f(q_1, \cdots, q_n, t) < 0$. Sometimes these are also included under 'non-holonomic.'

1.1 Holonomic constraints in disguise

Note that there are some special cases of *velocity-dependent* constraints which can actually be integrated to give holonomic constraints, these are holonomic-in-disguise. For example, consider a general velocity-dependent constraint:

$$A\dot{x} + B = 0.$$

This doesn't look holonomic, and in general it's not. However, in the case where there exists a function f such that $A = \partial f / \partial x$ and $B = \partial f / \partial t$, then

$$\begin{aligned} \frac{\partial f}{\partial x}\frac{dx}{dt} + \frac{\partial f}{\partial t} &= 0 \\ df &= 0 \end{aligned}$$

which means we may integrate this to yield a holonomic constraint $f - (\text{constant}) = 0$.

1.2 Non-holonomic constraints as general-ish constraints

In this sense we can always disguise a holonomic constraint as a non-holonomic constraint. Given $f(q,t) = 0$, just take the time derivative of this constraint and obtain a constraint which depends on $\dot{q}$ as well as q. Thus we can think of holonomic constraints as a special case of non-holonomic constraints; those which are integrable.

Modulo constraints that cannot be described by equations, we can then write constraints generically in terms of a set of equations $f(q, \dot{q}) = 0$. For most cases of interest (including all holonomic constraints), we can focus on constraints that are *linear* in $\dot{q}$:

$$\sum_{i,j} f_{ij}(q)\dot{q_j} = 0$$

1.3 Degrees of freedom

Hand & Finch describe non-integrable constraints (the particular type of non-holonomic constraint that we've focused on) in terms of a set of N_C equations

$$\omega_{ij}(q_1, \cdots, q_{N_C})\dot{q}_j = 0,$$

where N_C is greater than the number of degrees of freedom N of the system. In other words, the system is embedded into a higher-dimensional configuration space. This should feel very similar to Lagrange multipliers—you can add Lagrange multiplier terms to the Lagrangian as auxiliary (fake) degrees of freeom with no kinetic term. If we take $L \to L + \lambda f(q)$, the variation with respect to λ simply gives the constraint.

There are subtleties in counting degrees of freedom with non-holonomic system. For practical purposes, when you're counting degrees of freedom, use the following algorithm:

1. Count all of the pieces of information required to specify the system.
2. Subtract from this the number of equations that correlate any two degrees of freedom.

The result is the number of degrees of freedom.

For example, recall the last question of homework #1: a wheel rolling inside the circumference of a circular hole. There are two pieces of information: the angle θ corresponding to the center of the wheel and the angle φ that specifies the orientation of the wheel. The two angles are constrained by the requirement that the wheel rolls without slipping: the distance rolled by the wheel $\pi r_{\text{wheel}}\varphi$ is equal to the distance along the hole's circumference $\pi R_{\text{hole}}\theta$.

For non-holonomic constraints, one may require additional degrees of freedom to solve the system—these additional degrees of freedom are there to take into account the 'path dependence' of the configurations.

2 Properties of non-holonomic constraints

2.1 An example: unicycle

We discussed the penny rolling down an inclined plane as a prototype example of a non-holonomic constraint. A simpler example of a non-holonomic constraint (from Leinaas) is the motion of a unicyclist.

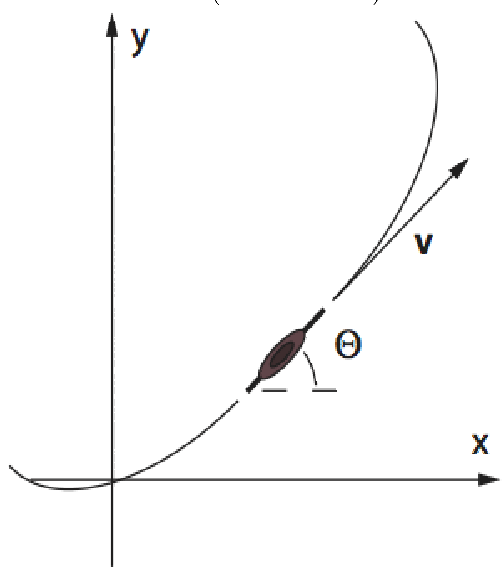

The position of the unicyclist is given by a pair of coordinates (x, y). An additional coordinate is the orientation of the unicycle, which is specified by an angle θ above. For any position (x, y) the unicyclist can be pointed in any direction, θ. Thus there is no constraint of the form $f(x, y, \theta, t) = 0$. Indeed, there are **three degrees of freedom**. However, we know that the unicyclist moves (non-slipping) only in the direction of θ. This means that there is a constraint that:

$$\dot{\boldsymbol{x}} = v(\cos\theta\,\hat{x} + \sin\theta\,\hat{y}),$$

or $\dot{y} = \dot{x}\tan\theta$. This is a non-holonomic constraint and so one cannot straightforwardly reduce the number of variables in the system.

2.2 Our definition of non-holonomic

Non-holonomic constraints are thos for which one must extremize a function $f(x,\ y,\ \cdots)$ such that the variation $\delta\boldsymbol{r}$ is perpendicular to some number of constraint vector fields $\vec{u}_i$. This is equivalent to a statement in terms of Lagrange multipliers; solve:

$$\vec{\nabla} f \;=\; \lambda_1\vec{u}_1 + \lambda_2\vec{u}_2 + \cdots$$

The case of a 'holonomic constraint in disguise' is when $\vec{u} = \vec{\nabla} g$. For example, in the case of one vector field, we have $\vec{\nabla} f = \vec{\nabla} g$ from which we can write $f - g = \text{const}$.

2.3 The non-holonomic constraint of the unicycle, explicitly

The unicycle configuration space has three dimensions: x, y, θ. We can write this as:

$$\frac{d}{dt}\begin{pmatrix} x \\ y \\ \theta \end{pmatrix} \;=\; \lambda_1\begin{pmatrix} 0 \\ 0 \\ 1 \end{pmatrix} + \lambda_2\begin{pmatrix} \cos\theta \\ \sin\theta \\ 0 \end{pmatrix}.$$

λ_1 tells us how the unicyclist is steering (changing θ), while λ_2 gives the forward motion where 'forward' is the direction that θ is pointing.

3 Holonomy and integrability

This is not a geometry class, so I'll leave out technical details. For those who want to see this properly, I encourage you to read Arnold's *Mathematical Methods of Classical Mechanics*, or books on general relativity. (The mathematically inclined should read up on differential geometry—there are several good books that connect physics and geometry[1].)

3.1 Geometry

Geometrically, **holonomy** refers to how an object transforms as it is "parallel transported" in a closed loop. For example, if one takes a vector tangent to the Earth at the equator, we have a notion of how to drag it in a way that maintains direction. Here's a picture from Wikipedia (under 'parallel transport')"

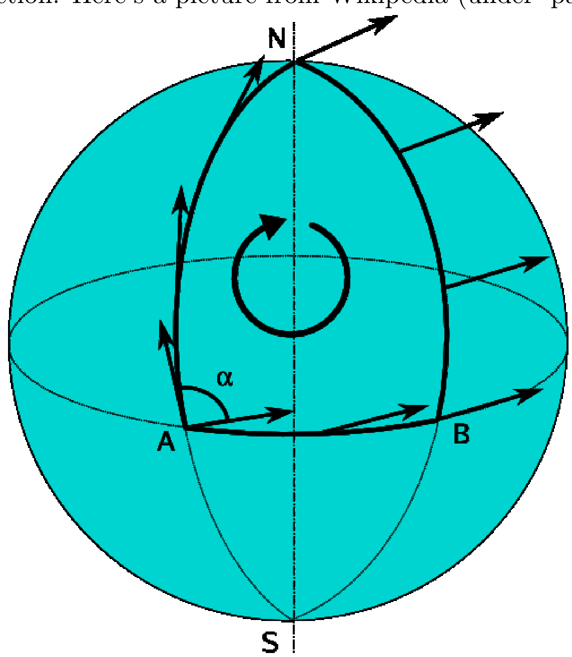

1. Off the top of my head: the graduate text *Geometry, Topology, and Physics* by Nakahara is a fantastic (albeit a bit advanced) reference. The huge tome by Misner, Thorne, and Wheeler is idiosyncratic but starts with the basics and should be accessible to you with careful study. I'm also especially fond of the book by Frankel in the references. The early portions of the review by Collinucci and Wijns (hep-th/0611201) also touches on the key topics.

Observe that the orientation of the vector changes after going around this closed loop! The geometry of the sphere tells us how much it changes, but the point is that the angel α is in general non-zero. **Holonomy** refers to how the vector can change after being pulled around 'parallel-ly' in a closed loop.

3.2 Non-integrability

This geometric notion of holonomy is related to non-integrability. The vector is being dragged around 'parallel-ly' so that we don't expect it to be changing. However, the geometry of the surface forces the vector to none-the-less change when it returns to its starting point.

In fact, the value of α depends on the particular closed path that we took. Compare this to high-school calculus: the integral of a total derivative only depends on its endpoints:

$$\int_{x_1}^{x_2} f(x)\,dx = \int_{x=x_1}^{x=x_2} dF(x) = F(x_2) - F(x_1).$$

Because the value of the holnomy α depends on the path (and not just the endpoints), it is in a sense **non-integrable**. In more familiar physics terminology, when the system is non-conservative we know that integrals over paths depend on the path and not just the endpoint.

3.3 (non)-holonomic constraints

Hertz first used the word holonomic (as well as non-holonomic/anholonomic) in the context of mechanical constraints. The point is that a constraint is **holonomic** if it can be integrated (to give $f(q,\,t)=0$) and then applied to reduce the degrees of freedom in a system. We should note that this is somewhat unfortunate nomenclature, since a **holonomic constraint** is one where there is no geometric holonomy (path-dependence), whereas a **non-holonomic constraint** is one that is related to this notion of geometric holonomy.

What do our (non)-holonomic *constraints* have to do with the idea of geometric holonomy? The whole point of classical dynamics is to show how a system changes in time; in other words, how does a point on the configuration space (e.g. (x, y, θ) for the unicycle) change once we give initial conditions $\left(x_0, y_0, \theta_0, \dot{x}_0, \dot{y}_0, \dot{\theta}_0\right)$? For a system with non-holonomic constraints, the state after some time evolution depends on the particular path taken to reach it. In other words, one can return to the original point in configuration space (x_0, y_0, θ_0) but *not* return to the original state.

For an explicit example of this that is related to the notion of parallel transport around the globe, see the 'nonholonomy o the rolling sphere' by Johnson in American Mathematical Monthly (reference below).

4 Geometry

In past discussion sections we've introduced the notions of a manifold M and its tangent space at a point q, T_qM. In summary, a point on space of possible configurations M is given by a vector q. The [generalized] velocity of a particular configuration is given by $\dot{q}$, an element of T_qM. The space of configurations combined with the possible velocities at each point is called the tangent bundle, TM. For our purposes you can think of $TM = M \times T_qM$ for some representative q. Recall that the Lagrangian L is a function from TM to the real numbers, $\mathbb{R}$.

4.1 Vector fields

A **vector field** is a function from $q \in M \to T_qM$ that assigns a tangent vector to each point in the manifold. That is, at every point in configuration space, you assign some particular velocity. For 'nice' (e.g. holonomic) physical systems, the goal of all of classical dynamics is to be able to draw the vector field representing the dynamical evolution of a system at each configuration.

The picture is that you have a vector field over a space M spanned by $(q, \dot{q})$. At $t=0$ you start at some point representing an initial configuration. The equations of motion give a vector field in phase space. At each time increment, you follow the vector to the next point in configuration space. In this way you trace out a path in phase space: you've defined a curve whose tangent vector is everywhere parallel to a given vector field. This is a *flow* along the vector field.

4.2 Distributions are constraints

A **distribution** is a map from each point q of manifold M to some subspace of T_qM. For example, in the case of the unicycle, the non-holonomic constraint is a distribution that restricts the generalized velocity to lie on a subset of allowed velocities. Recall:

$$\frac{d}{dt}\begin{pmatrix} x \\ y \\ \theta \end{pmatrix} = \begin{pmatrix} \dot{x} \\ \dot{y} \\ \dot{\theta} \end{pmatrix} = \lambda_1 \begin{pmatrix} 0 \\ 0 \\ 1 \end{pmatrix} + \lambda_2 \begin{pmatrix} \cos\theta \\ \sin\theta \\ 0 \end{pmatrix}.$$

Physically the λ_2 is saying that the velocity of the unicycle on the plane must be pointing in the direction specified by θ. Meanwhile the λ_1 term says that θ can vary independently of x and y. This equation defines a subspace of the three dimensional tanget space of the system.

A **regular distribution** is constant over the configuration space.

4.3 The Parallel Parking Product

If you're familiar with differential geometry, you'll know about the **Lie bracket**. If not, then any attempt to explain it to you here will turn you off. Instead, let me offer Mason's suggestion that this should be called the 'parallel parking product.'

The idea is this. Suppose you have two vector fields, v and w. The **parallel parking product** (or **Lie bracket**) is a way to measure the noncommutativity of flows along these vector fields. It's written as $[v, w]$, and it's no coincidence that it looks like the commutator that appears in quantum mechanics courses. Consider the following image from the web[2]:

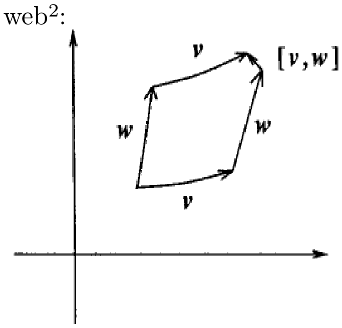

Fig. 10. $[v, w] \neq 0$

The value of $[v, w]$ is what happens when you start at the bottom left corner and compare the following two end points:

1. First flow along v for a short time, then flow along w for a short time.
2. First flow along w for a short time, then flow along v for a short time.

The difference in these two flows as 'short time' becomes infinitely short is the Lie bracket $[v, w]$. This can also be thought of as a measure of how close you are to closing a path by flowing along v then w then $-v$ and then $-w$. Recall that we made a big deal about how things change along closed paths when we introduced the geometric notion of holonomy.

Note that the output of a Lie bracket is a vector at each point in time, so the Lie bracket is itself a vector field.

2. `http://math.stackexchange.com/questions/163262/visualizing-commutator-of-two-vector-fields`

4.4 Involutions are holonomic constraints

Recall that a **distribution** is basically the manifestation of our non-holonomic constraint: it is a restriction of the tangent space (velocity space); i.e. a restriction of velocities. The distribution is then spanned by some number of vector fields; i.e. at each point in space q the distribution assigns a subspace of T_qM which is spanned by some vectors.

Each Lie brackets between any pair of these vector fields is a new vector field. If, for every such Lie bracket is contained in the distribution (i.e. is a linear combination of the vector fields we started with), then we say that the distribution is **involutive** or is an **involution**.

In other words, an involution is a distribution which is closed under the Lie bracket.

We will relate this to holonomic constraints below. (Recall from above that all holonomic constraints can be written as seemingly-non-holonomic constraints that are integrable!)

4.5 The Frobenius theorem and parallel parking

The problem of parallel parking (or the inverse: getting out of a tight parking spot) is completely analogous to the motion of the unicycle: you can steer left or right θ but you can only move forward or backward for a given θ. How to you make a sideways motion? You have to shimmy in by going back and forth. This shimmying is captured by the Lie bracket.

For parallel parkers this is great: naively your only allowed motions are along the vector fields defined by θ and 'forward/backward.' But by taking a Lie bracket (i.e. forward, turn, backward, turn the other way) you end up in a place that's not quite the same as the original orientation... you can make a small sideways motion.

The punchline for all this is that there's a theorem in geometry called the **Frobenius theorem**. It says that a regular distribution is **integrable if and only if it is involutive**. What does this all mean for us? Mason summarizes Frobenius theorem as follows: if a system as nonholonomic constraints, then parallel parking is useful. Mason, it should be noted, is interested in robotics and trying to get the most out of a small number of robitic motions.

For us, the Frobenius theorem tells us something else. We *want* integrable flows: these are the solutions to our equations of motion: the path in configuration space as a function of time that describe the dynamical evolution of the system. The Frobenius theorem tells that we can only construct such paths if the constraints are involutive—that is, *involutions correspond to holonomic constraints*. Involutions allow us to transfer the constraints on the tangent bundle TM to the space M by virtue of integrating the paths.

Non-holonomic constraints are given by distributions which are not involutive.

4.6 Application to the unicycle

For a technical demonstration, we just have to write out the Lie bracket more technically. For vector fields f_i and g_i, define the Lie bracket to be

$$[f, g]_i = \frac{\partial g_i}{\partial q_j} f_j - \frac{\partial f_i}{\partial q_j} g_j.$$

You can take this as a definition or otherwise think carefully about why this represents the conceptual definition above. *Hint*: the partial derivative $\partial g_i / \partial q_j$ is how much the vector field g_i changes as I move along some direction q_j. Dotting this into f_j means the change in g_i as I move infinitesimally along the flow generated by f_j.

In the unicycle (or parking) problem, the two vector fields are those generated by the allowed directions in the tangent space (the allowed velocities):

$$g = \begin{pmatrix} 0 \\ 0 \\ 1 \end{pmatrix}$$
$$f = \begin{pmatrix} \cos\theta \\ \sin\theta \\ 0 \end{pmatrix}.$$

Then the Lie brackets are:

$$\begin{aligned}
[g, f] &= \begin{pmatrix} 0 & 0 & 0 \\ 0 & 0 & 0 \\ 0 & 0 & 0 \end{pmatrix} \begin{pmatrix} \cos\theta \\ \sin\theta \\ 0 \end{pmatrix} - \begin{pmatrix} 0 & 0 & -\sin\theta \\ 0 & 0 & \cos\theta \\ 0 & 0 & 0 \end{pmatrix} \begin{pmatrix} 0 \\ 0 \\ 1 \end{pmatrix} \\
&= \begin{pmatrix} -\sin\theta \\ \cos\theta \\ 0 \end{pmatrix}.
\end{aligned}$$

This is not in the span of $\{g, f\}$ and hence the system is nonholonomic/not involutive/not integrable, etc. etc. etc.

4.7 The holonomic limit of the unicycle

Suppose that instead of a unicycle, the system described a car with a broken steering wheel. Mason shows that this is now a holonomic system. This changes the constraints:

$$\begin{aligned}
\begin{pmatrix} 0 \\ 0 \\ 1 \end{pmatrix} \cdot \begin{pmatrix} \dot{x} \\ \dot{y} \\ \dot{\theta} \end{pmatrix} &= 0 \\
\begin{pmatrix} \sin\theta \\ -\cos\theta \\ 0 \end{pmatrix} \cdot \begin{pmatrix} \dot{x} \\ \dot{y} \\ \dot{\theta} \end{pmatrix} &= 0.
\end{aligned}$$

These equations say that (1) you can't turn and (2) that you can't move sideways (perpendicular to θ). Thus:

$$\begin{aligned}
\dot{x}\sin\theta - \dot{y}\cos\theta &= 0 \\
\dot{\theta} &= 0.
\end{aligned}$$

These equations can be integrated:

$$\begin{aligned}
\theta &= \theta_0 \\
(x - x_0)\sin\theta_0 - (y - y_0)\cos\theta_0 &= 0.
\end{aligned}$$

5 The rolling penny problem

The case of a rolling penny on a flat plane is dicussed in Hand & Finch appendix A. Please make a point to read this carefully. For more details about the problem of the rolling penny on an inclined plane that we discussed in class, see the American Journal of Physics article by Flannery cited below. It is a little technical but written at a level that you should now be able to undrerstand. Most of the document focuses on the question of whether Lagrange multipliers can be used for non-holonomic constraints—this is a very interesting question which is summareized at the end of section II.C of that document. The discussion of the solution of the inclined plane problem is given in section III. The point is that Lagrange multipliers can be used for non-holonomic constraints that are linear in $\dot{q}$; Flannery shows how to map the 'penny on an inclined plane' constraints to linear constraints.

6 Nonholonomic constraints and Lagrange multipliers

This is based mostly on Flannery and the summary by Wolf (who in turn summarizes the textbook by Bloch), references below. You know that a nice way to include a holonomic constraint $f(q) = 0$ into the Lagrangian formalism is to use Lagrange multipliers:

$$L' = L + \lambda(t) f(q).$$

If we pretend that the Lagrange multiplier λ is an auxilliary (*fake*) degree of freedom, its equation of motion enforces the constraint. The equation of motion of the actual degrees of freedom will now have the constraint folded in. One solves for the Lagrange multiplier λ (proportional to the restoring force) and plugs into the solution for the actual degrees of freedom.

In order to solve problems with non-holonomic constraints, one might apply the same trick. Given a non-holonomic constraint $g(q,\dot{q})=0$, can we can try to write

$$L' \;=\; L + \lambda(t) g(q,\dot{q}).$$

There's a huge subtlety here—see the erratum to Ray's article in the American Journal of Physics. It turns out to be related to whether the path variations $\delta q(t)$ are required to satisfy the non-holonomic constraints or whether those constraints are imposed after taking variations. If you do this incorrectly, you run afoul of what you would derive from the Newtonian formalism.

The punchline is that for non-holonomic constraints that are *linear* in $\dot{q}$ the method of Lagrange multipliers works (although their validity is rather subtle!).

7 Other things I'd like to get to (but may not have time for)

- The falling cat problem (see Wikipedia article for a cute graphic and a list of reference). Interesting remark: some of the early work on this type of problem was done by Frank Wilczek, who applied techniques related to the strong nuclear force which later earned him the Nobel Prize in Physics.
- The parallel parking problem (again, see Wikipedia article for references or google). Turns out to be related to the falling cat problem. For some details, see page 33 of
  `https://web.math.princeton.edu/~nelson/books/ta.pdf`
- Engineers seem to care about non-holonomic constraints when building robots with the fewest number of joints required to achieve a range of motion. See Mason's slides in the reference below.

8 References

- Hand & Finch, *Analytical Mechanics*
- Frankel, *The Geometry of Physics*
- Jose & Saletan, *Classical Dynamics: A contemporary approach*
- Johns, *Analytical Mechanics for Relavitity and Quantum Mechanics*
- Bloch, Marsden, Zenkov. "Nonholonomic Dynamics" in Notices of the AMS, Vol 52, 324.
- Jon Leinaas, Classical Mechanics and Electrodynamics, FYS 3120 Lecture Notes (2009)
- Berry, "Anticipations of the Geometric Phase" in Physics Today, December 1990
- B.D. Johnson, "The Nonholonomy of the rolling sphere," American Mathematical Monthly, July-July 2007, Vol 114, 500. `http://euler.slu.edu/~johnson/public/maths/bdjsphere.pdf`
- Mason, Mechanics of Manipulation (CMU), Spring 2013, lecture 5 slides available online: `http://www.cs.cmu.edu/afs/cs/academic/class/16741-s07/www/lecture5.pdf`
- Flannery, "The enigma of nonholonomic constraints," Am. J. Phys. **73** (2006) 265.
- Bloch, *Nonholonomic Mechanics and Control*.
- Michael Wolf's CDS 205 (Caltech, 2005) final project. Available online: `http://www.cds.caltech.edu/~marsden/wiki/uploads/projects/geomech/Wolf2005.pdf`
- Ray, "Nonholonomic Constraints," Am. J. Phys. **34** (1966) 406 and erratum (page 1202)