Control theory might be different than you think it is

Andrew D. Lewis

Department of Mathematics and Statistics
Queen's University, Kingston, ON, Canada

2016/05/09 (CDSR 2016)

What we *are not* talking about

- We are not devising design methodologies useful for solving specific problems or restricted classes of problems.
- We are not talking about control strategies, e.g., output regulation, backstepping, passivity-based control, L^1-adaptive control, sliding mode control, etc.
- Not mentioning these things does not equate to our belief that they are not worth talking about. There are many things we find important and/or interesting that we are not going to talk about.
- But our objectives are different.

What we *are* talking about

- We are interested in understanding fundamental system structure.
- This structure pertains to problems such as controllability, stabilisability, and optimality.
- These problems have been studied in depth and detail since at least the 1960's... nonetheless, the understanding of these aspects of system structure is very far from complete.
- These problems are not easy—to make progress with them one has to understand very well what one is doing.

Collaborators and students

- Over the years: César Aguilar, Francesco Bullo, Bahman Gharesifard, Ron Hirschorn, Saber Jafarpour, Abdol-Reza Mansouri, Richard Murray, David Tyner.

The arc of the talk

- The end objective is to provide a list of requirements of a correct and useful[1] mathematical theory of control.
- We shall do this by a sequence of illustrations—some quite concrete, some less so—that point out the importance of certain ideas.
- Many of our illustrations will have a mechanical flavour, but the end result is a *general* theory of control.

[1] I am a mathematician, remember, so "useful" means proving theorems.

Some mechanical control systems

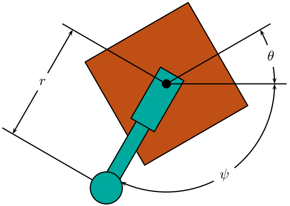

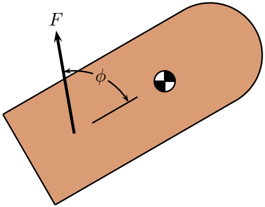

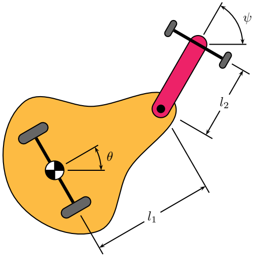

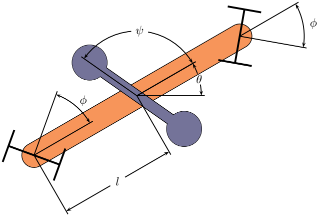

Snakeboard gait: x Snakeboard gait: y Snakeboard gait: θ

Motion planning and controllability

- The natural problem for such systems is *motion planning:* Steer the system from one state to another.
- The associated existential problem is *controllability:* Can we steer the system from one state to another?
- There is much work on the motion planning problem in the robotics literature, some of it for underactuated systems such as we are interested in.
- And yet... the controllability problem is utterly unresolved in any practical degree of generality.
- Let us illustrate how one might come across this limitation in practice.

Illustrating the controllability problem

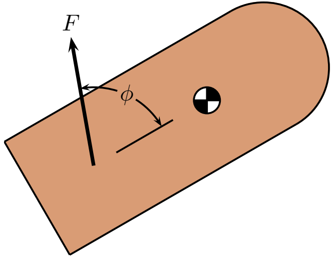

- Hovercraft system:
  1. *Question:* Is the system controllable?
  2. *Answer:* Yes, and can be proved using "standard results.[2]"

[2] e.g., Sussmann, *SIAM J. Control Optim.*, **25**(1), 158–194, 1987

Illustrating the controllability problem (cont'd)

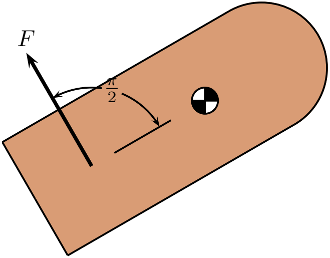

- Now suppose that the fan cannot rotate.
  1. *Question:* Is the system controllable?
  2. *Answer:* No, but now requires a special theorem.[3]

[3] Bullo/L, *SIAM J. Control Optim.*, **44**(3), 885–908, 2005

A few words about the mathematical approach to these problems

- The controllability for the preceding examples is proven using "geometric control of mechanical systems," as developed by Bullo/L.[4]
- Mechanical systems are converted into a differential geometric model known as an "affine connection control system," and tools from affine differential geometry are used to prove controllability theorems.
- The moral of the mathematical modelling story here is this: *Thou shalt only work with objects that are independent of choices of coordinates.*

---

[4] Springer-Verlag, *Texts in Applied Mathematics*, 2004

A few words about the mathematical approach to these problems (cont'd)

- Here is something that *is not* independent of coordinates. The equations of motion for a robotic system take the general form

$$\underbrace{\overbrace{\overbrace{M(q(t))\ddot{q}(t)}^{\text{“inertial term”}}+\overbrace{C(q(t),\dot{q}(t))}^{\text{“Coriolis term”}}}^{\text{Individually not well-defined!!}}}_{\text{These must be kept together!!}}=F(t,q(t),\dot{q}(t)).$$

- Interestingly, the mathematics leads to answers to the motion planning problem as well!

Hovercraft: Decouple1 Decouple2 Plan1 Plan2 Plan3

Snakeboard: Plan1 Plan2

Summary of where we are

- There is a very nice “body of work” here, with interplay between applications and interesting mathematics.
- The mathematics is at a level that a course in the material can be taught to 4th-year undergraduates.
- And yet…the fundamental existential question remains unanswered…
- *Punchline(s)*
  1. Even with simple problems, one can easily run up against the limits of what the known theory can tell you.
  2. Sometimes thinking about an applied problem in a mathematical way can help solve problems that may be difficult or impossible to solve otherwise.

Some more mechanical control systems

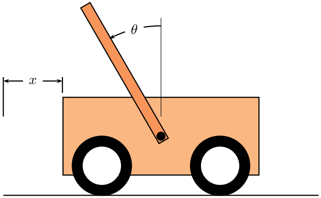

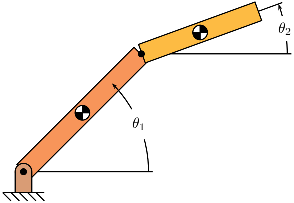

- The problem here is stabilisation of the unstable equilibria.
- These problems can be dealt with locally by using linearisation and linear control theory.
- Can one do better by using a nonlinear stabilisation methodology?

Stabilisation and stabilisability

- For mechanical systems, stability is related to properties of the potential energy function:

- *Problem:* Using feedback, can we turn a mechanical system with an unstable equilibrium into a mechanical system with the same equilibrium, now stable?
- This is an easy to understand and physically natural problem.

Stabilisation and stabilisability

- There's a lot of work here, under the general name of "energy shaping:"
  - Takegaki/Arimoto, *Trans. ASME Ser. G*, **103**(2), 119–125, 1981
  - van der Schaft, *Nonlinear Anal. TMA*, **10**(10), 1021–1035, 1986
  - Ortega, et al. *IEEE Trans. Automat. Control*, **47**(8), 1218–1233, 2002
  - Bloch, et al. *IEEE Trans. Automat. Control*, **45**(12), 2000, 2253–2270 and **46**(10), 1556–1571, 2001
  - Blankenstein, et al. *Int. J. Control*, **75**(9), 645–665, 2002
  - Chang, et al. *ESAIM Control Optim. Calc. Var.*, **8**, 393–422, 2002
  - etc., etc., etc.
- *Existential question:* Is it possible to determine when the energy shaping idea works?
- *Broader existential question:* Is it possible to recognise when a control system is stabilisable?

Feasibility of energy shaping

- The method of energy shaping leads to overdetermined quasilinear partial differential equations.
- The mere *existence* of solutions to such equations is highly problematic, never mind existence of solutions with properties such as "stabilising."
- Despite these deep existential problems, the method has been pursued.
- Tiny bit of work on the existence problem:
  - Auckly/Kapitanski, *SIAM J. Control Optim.*, **41**(5), 1372–1388
  - Gharesifard, et al. *Commun. Inf. Syst.*, **8**(4), 353–398
- *Punchline:* The matter of determining, *in any generality*, when the energy shaping method can be applied seems hopeless.

A few words about the mathematical approach to these problems

- This is really not trivial: e.g., it involves enough homological algebra to make sense of the following diagram

$$
\begin{tikzcd}
 & & 0 \arrow[d] & 0 \arrow[d] & & \\
 & & \mathsf{S}^2(\mathsf{T}^*\mathsf{M}) \arrow[r, "\sigma_1(\Phi)"] \arrow[d] & \mathsf{T}^*\mathsf{M} \otimes \mathcal{F} \arrow[r, "\tau"] \arrow[d] & \mathsf{K} \arrow[r] & 0 \\
0 \arrow[r] & \rho_1(\mathsf{R}) \arrow[r] \arrow[d] & \mathsf{J}^2\mathsf{E} \arrow[r, "\rho_1(\Phi)"] \arrow[d] & \mathsf{J}^1\mathcal{F} \arrow[d] & & \\
0 \arrow[r] & \mathsf{R} \arrow[r] & \mathsf{J}^1\mathsf{E} \arrow[r, "\Phi"] \arrow[d] & \mathcal{F} \arrow[d] & & \\
 & & 0 & 0 & &
\end{tikzcd}
$$

which is used to construct the compatibility operator as a map from the bottom left corner to the top right corner.

Summary of where we are

- With energy shaping and stabilisability, the picture is *a lot* different than with motion planning and controllability.
- There is not really a closed "body of work" here. There are isolated examples and a theory that seems impossible to work with.
- There are interesting *mathematical* problems here, and these require deep, difficult, and specialised knowledge.
- What "stabilisation and stabilisability" have in common with "motion planning and controllability" is this: the fundamental existential questions are entirely unanswered.
- *Punchline(s)*
  1. Sometimes natural physical problems give rise to very difficult mathematical problems.
  2. Sometimes thinking about an applied problem in a mathematical way can reveal that the approach is simply not a good one.

Yet another mechanical control system

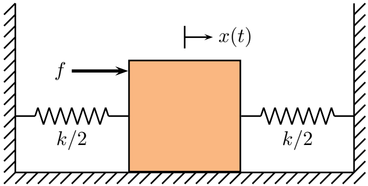

- Force is bounded: $f \in [-A, A]$.
- *Problem:* Steer from $(x(0), \dot{x}(0))$ to $(0, 0)$ in time T while minimising

$$\int_0^T \frac{1}{2} k x^2(t)\, \mathrm{d}t.$$

- Another simple physical problem.

Measurable controls are required

- There is a solution to this optimisation problem, and the optimal solution involves the force f "chattering" infinitely often as the final state is approached.
- This is the ***Fuller phenomenon***.[5]
- Therefore, continuous, or piecewise continuous, controls are not sufficient in any general theory of control.
- The correct class of controls to use is Lebesgue measurable controls.

[5] *IFAC World Congress*, Moscow, 1960

Physical models are real analytic

- A typical control system, in a general formulation, has the form

$$\dot{x}(t) = f(x(t), u(t)),$$

  with x being the state, u being the control, and f prescribing the dynamics.
- What properties should f have? In particular, what should be the nature of the dependence of f on the state x?
- Typical modelling frameworks have f being smooth (infinitely differentiable) with respect to x.
- This has the advantage of being mathematically convenient.
- Physically, it is quite the wrong assumption. Physical systems are either
  1. nonsmooth or
  2. real analytic.

Why real analyticity is important

- Important results are true for real analytic systems that are not true for smooth systems:
  1. usefully general versions of Frobenius's Theorem;
  2. sufficiency of Lie brackets for determining certain motion planning problems;
  3. usefully general versions of the Orbit Theorem.
- One often sees statements about conditions holding on "open dense sets," e.g., a controller is globally stabilising except on an open dense set. What is actually true in most of these cases is that the set has an analytic complement, and this is a *much* stronger and more useful statement.
- Real analytic things behave like you think smooth things should.

Why real analyticity is hard

- There are no real analytic partitions of unity.
- The things that one can achieve for smooth objects using partitions of unity can often also be achieved for real analytic objects. . . using sheaf cohomology. . . gulp. . .
- Appropriate notions of "convergence of a sequence of real analytic functions" are difficult to define.
- After a lot of work. . . these notions of convergence *can* be defined in ways that *can* be worked with.

Summary of where we are

- A comprehensive and physically realistic framework for control theory *must* incorporate measurable controls and really *should* incorporate real analytic system models.
- *Punchline(s)*
  1. A really honest framework for control theory must be done using some difficult (and sort of standard) methods, and some really difficult (and definitely nonstandard) methods.

What is the right model for a control system?

- We have already mentioned the models of the form
$$\dot{x}(t) = f(x(t), u(t)).$$
- We have already mentioned the importance of working with constructions that do not depend on particular choices of coordinates.
- What about the dependence of the dynamics on control? Is there a sort of analogue of "coordinate independence" for this?
- There is. . . and you're not going to like it. . .

Linearisation is not what you think it is

- First let's think about linearising dynamics.
  - Have a differential equation $\dot{x}(t) = f(x(t))$.
  - Have an equilibrium point x_0, i.e., $f(x_0) = 0$.
  - Linearise: define $A = \frac{\partial f}{\partial x}(x_0)$.
- Good news! This makes sense and one has results like

$$\dot{x}(t) = Ax(t) \text{ asymptotically stable}$$
$$\Downarrow$$
$$\dot{x}(t) = f(x(t)) \text{ locally asymptotically stable.}$$

Linearisation is not what you think it is (cont'd)

- Now consider linearising a control system.
  - Have a control system $\dot{x}(t) = f(x(t), u(t))$.
  - Have an equilibrium point (x_0, u_0), i.e., $f(x_0, u_0) = 0$.
  - Linearise: define $A = \frac{\partial f}{\partial x}(x_0, u_0)$ and $B = \frac{\partial f}{\partial u}(x_0, u_0)$.
- Some good news: one has results like

$$\dot{x}(t) = Ax(t) + Bu(t) \text{ controllable (stabilisable)}$$

$$\Downarrow$$

$$\dot{x}(t) = f(x(t), u(t)) \text{ locally controllable (stabilisable).}$$

- But there is bad news: the process does not make sense.

Linearisation is not what you think it is (cont'd)

- Consider two systems:

$$\begin{aligned} \dot{x}_1(t) &= x_2(t), & \dot{x}_1(t) &= x_2(t), \\ \dot{x}_2(t) &= x_3(t)u_1(t), & \dot{x}_2(t) &= x_3(t) + x_3(t)u_1(t), \\ \dot{x}_3(t) &= u_2(t), & \dot{x}_3(t) &= u_2(t), \end{aligned}$$

  with the equilibrium point $((0,0,0),(0,0))$.
- These systems have *exactly* the same trajectories and so are "the same."
- The linearisation of the left system is neither controllable nor stabilisable and the linearisation of the right system is controllable (and so stabilisable).
- *Questions:* What is the linearisation of a system? What does it mean for a system to be linearly controllable?

Summary of where we are

- The problem with the preceding example is this: The two different systems are really the same system, but the parameterisation of the dynamics by control is different.
- Just like one wants "coordinate invariance," one also wants "control parameterisation invariance."
- *Punchline(s)*
  1. There are lots of common constructions in control theory that, while (hopefully) independent of coordinates, are not independent of control parameterisation.
  2. One would like to eliminate the dependence of models on control parameterisation.

Control sets should be general

- In a model like $\dot{x}(t) = f(x(t), u(t))$, the set in which the control u lives can be quite general, e.g., discrete, a polyhedron, etc.
- Control sets should *never* be assumed to be smooth, e.g., if you have a framework where you will be differentiating with respect to control. . . this is done *with* loss of generality. . .
- Smooth structure is often assumed for control sets:
    1. control sets are sometimes assumed to be open, e.g., in "dynamic feedback linearisation" and "differential flatness";
    2. there is the "bundle version" of a nonlinear system, where control sets are assumed smooth.[6,7]
- Assumptions of smooth structure are mathematically convenient, but ultimately lacking in practicality.

---

[6] Brockett, in *The 1976 Ames Research Center (NASA) Conference*, 1–14, 1976

[7] Willems, *Ricerche Automat.*, **10**(2), 71–106, 1979

Systems are only locally defined

- There are many reasons for working with systems where data is only locally defined.
- There are application oriented reasons, e.g.,

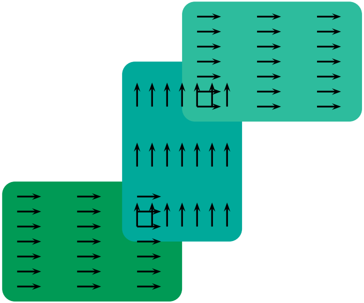

- There are mathematical reasons, e.g., solutions to (globally defined) differential equations may only be locally defined.

A summary of all of summaries of where we are

- We have arrived at the following list of criteria for a useful kind of control theory:
    1. Must be coordinate-invariant
    2. Must be control parameterisation-invariant
    3. Must work with general control sets
    4. Must make possible measurable controls
    5. Should seamlessly incorporate real analytic models
    6. May permit locally defined data
- Developing a framework for doing all of this is difficult. You need:
    1. Differential geometry
    2. Advanced functional analysis (locally convex topologies)
    3. A theory of measure and integration in infinite-dimensional spaces
    4. Sheaves and groupoids

References

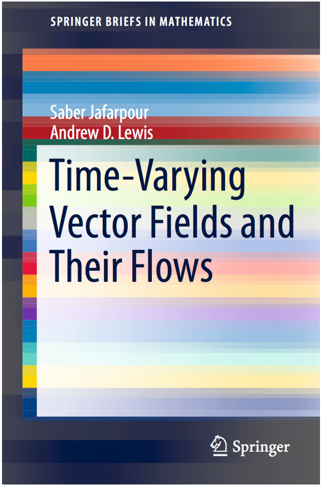
SPRINGER BRIEFS IN MATHEMATICS
Saber Jafarpour
Andrew D. Lewis
Time-Varying Vector Fields and Their Flows
Springer

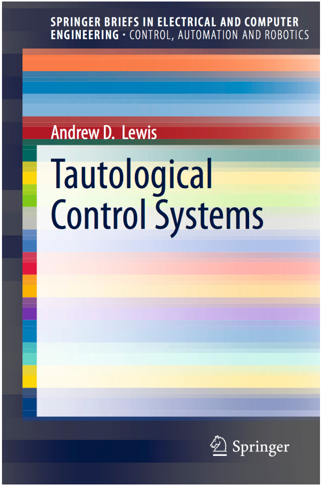
SPRINGER BRIEFS IN ELECTRICAL AND COMPUTER ENGINEERING · CONTROL, AUTOMATION AND ROBOTICS
Andrew D. Lewis
Tautological Control Systems
Springer

Take-away message

- Some problems, even concrete ones, in control theory can benefit from a little mathematics.
- Not all problems in control theory can benefit from a little mathematics.
- The basic structural problems in control theory remain open, and require some rethinking of the framework if they are to be answered.
- Understand what you are being told. . . *then question it!*