On The Teaching of Nonlinear Control: Challenges and Ideas

Manfredi Maggiore, *Senior Member, IEEE*

(Opinion Paper)

***Abstract*—This letter discusses the teaching of nonlinear control courses at the graduate level with emphasis on control design, presenting the views of the author about some of the challenges that the discipline is facing. The main concern raised in this letter is the current focus of the discipline on problems of equilibrium stabilization, both in teaching and research. The author contends that this focus is detrimental and it limits the perceived relevance of nonlinear control courses to graduate students working in other areas. The author proposes that changing the focus of graduate courses to problems of set stabilization may allow the teacher to better motivate nonlinear control concepts with exciting examples, while at the same time establishing connections with other disciplines. A sample lecture and assignment are presented illustrating the ideas put forward in this letter.**

***Index Terms*—Control engineering education, nonlinear control systems, robot control.**

I. Nonlinear Control in Context

NONLINEAR control (NLC) is a subject with a long history, going at least as far back as[1] Lyapunov's 1892 work on stability theory [3]. The first applications of nonlinear control might have occurred in the 1930s, when Četaev and other Russian scientists recognized the central importance of stability theory in guidance and control systems for aviation.

Modern NLC is highly interdisciplinary, being intertwined with a number of scientific areas, among which dynamical systems, mechanics, differential geometry and the calculus of variations have historically been important for its evolution. In turn, these areas have been greatly impacted by developments in NLC. In the 1970s and 80s, NLC developed a close relationship with the area of robotics (more on this later), and today NLC tools have become indispensable in a variety of disciplines: systems biology, distributed robotics, and aerospace engineering, to name a few.

As the area of NLC matured, the teaching of the subject evolved, and today one finds two typical course offerings at the graduate level. One type is centred on stability theory, and it covers the dynamical systems perspective of NLC. The second type is centred on differential geometry, and it covers structural aspects of NLC systems, such as controllability and feedback equivalence.

The most popular textbook for stability-oriented courses is Khalil's *Nonlinear Systems* [4]. Other popular references are [5], [6]. Popular textbooks for geometry-oriented courses are Isidori's *Nonlinear Control Systems* [7] and Nijmeijer and Van der Schaft's *Nonlinear Dynamical Control Systems* [8]. The book by Bullo and Lewis [9] is a popular reference for geometric control of mechanical systems. Other references include [10], [11].

II. The Emphasis on Equilibrium Stabilization

A conventional NLC course focused on control design[2] does not spend much time discussing control specifications. Often, the control objective is the stabilization of an equilibrium or the asymptotic tracking of a reference signal, this latter also cast as equilibrium stabilization for a time-varying control system.

The variety of control tools one will learn in the course does not arise from an assortment of control specifications, but instead from ways to stabilize the equilibrium: adaptively, robustly, discontinuously, via output-feedback, dynamically, event-triggered, network-based, and so on. The same emphasis on equilibrium stabilization dominates our textbooks and is a reflection of the emphasis that the field of NLC has placed on this problem.

While it might be clear to a student why it is relevant, say, to develop an adaptive controller for a system with uncertain parameters, it is far from obvious why the control specification should be stabilizing an equilibrium. Many research topics of recent interest require non-equilibrium tools. For example, typical problems in distributed robotics involve the stabilization of *sets* [12], and so do locomotion problems in robotics [13]. A student interested in these research areas will be puzzled by the emphasis on equilibria dominating our courses and our textbooks, and will unavoidably look elsewhere for tools.

If graduate students don't see the connections between NLC and their research area, NLC will become increasingly

Manuscript received 6 June 2022; revised 8 August 2022; accepted 18 August 2022. Date of publication 23 August 2022; date of current version 7 September 2022. This work was supported by the Natural Sciences and Engineering Research Council of Canada (NSERC). Recommended by Senior Editor C. Prieur.

The author is with the Department of Electrical and Computer Engineering, University of Toronto, Toronto, ON M5S 3G4, Canada (e-mail: maggiore@control.utoronto.ca).

Digital Object Identifier 10.1109/LCSYS.2022.3200913

[1] In [1], Loría and Panteley trace the origin of stability theory back to Lagrange's 1788 treatise *Méchanique Analitique* [2].

[2] There are other important topics of NLC that do not pertain control design. These are typically structural and concern properties such as controllability, observability, disturbance decoupling that are taught in geometric nonlinear control courses. The discussion in this section is focused on control design.

isolated. There might already be signs of distress in this discipline, judging from the number sessions dedicated to nonlinear control systems at the IEEE CDC.[3]

In this letter I will use the area of robotics to motivate my discussion, but by no means the observations I make are limited to this particular research field. I choose robotics simply because this area is close to my own research interests and I can formulate opinions based on direct experience. In the next Section I examine the historical connections between NLC and robotics, and draw some lessons from the current relationship between these two areas.

III. Nonlinear Control and Robotics

During the 1970s and 80s, the fields of nonlinear control and robotics witnessed a beneficial cross-fertilization, with developments in either one area having direct effects on the other. In this section, I describe an example of this beneficial cross-fertilization. Then I present my opinion about the status of this relationship, connecting it to the discussion in the previous section.

A Virtuous Loop (From Computed Torque to Biped Walking): The computed torque method is a technique developed in the early 1970s to control fully actuated robots by cancelling out nonlinearities and rendering the model linear. This idea, commonly attributed to the works in [14], [15], spread quickly in the robotics community, with a number of experimental results demonstrating its practical viability.

Shortly after the development of the computed torque method, Brockett [16, Th. 2] in 1978 formulated and solved, for single-input nonlinear control systems, a restricted version of the feedback linearization problem. In Brockett's words, the problem is to determine "Which systems are feedback equivalent to n-th order integrators?". Feedback linearization is a generalization of the computed torque method because in this latter the feedback transformation gives parallel double-integrators whose inputs are the joint torques and whose outputs are the robot's joint variables.

Brockett's paper [16] was followed by a number of papers in the early 1980s aimed at generalizing his results [17], [18], [19]. Eventually, this line of research led to the introduction by Byrnes and Isidori [20], [21] of the concept of zero dynamics, an idea closely related to input-output feedback linearization.

In the early 2000s, the zero dynamics, originally developed with a focus on the stabilization of equilibria [21], became a crucial ingredient for the design controllers inducing stable walking in biped robots [13], and today the keyword *hybrid zero dynamics* has come to represent an approach to robot locomotion.

The foregoing discussion illustrates how an idea originating in robotics migrated to the field of NLC stimulating important developments there, and eventually making its way back to robotics. This is only one of many examples of the cross-fertilization of robotics and NLC.

Divergence of the Two Fields: The relationship between robotics and NLC appears to be waning. It is rare today to find NLC ideas making an impact in robotics, with the important exception of the work in [13] just mentioned.

Much focus in modern robotics is on vision, machine learning, localization, and so on, but many robotics problems would benefit from NLC tools. For instance, how does one make a humanoid robot climb stairs, sit on a chair, and stand up, with proven stability? These research questions fall firmly within the purview of NLC.

The disconnect between these two areas is also occurring on the educational front. In my institution, while twenty years ago graduate students specializing in robotics would take an NLC course, today they are more likely to take courses in optimization and computer science. It is hardly surprising that NLC is not on their radar. A graduate student wanting to make a robot climb stairs will not see how stabilizing an equilibrium can help them solve the problem, and will turn elsewhere for answers.

Criticism: The situation described above represents a missed opportunity for both disciplines. Modern problems of robotics could spur new and exciting theoretical developments in NLC. At the same time, robotics would benefit from algorithms based on solid grounds, as this discipline increasingly relies on heuristic techniques that can't be assessed on objective grounds.

We NLC theorists need to place greater emphasis on control specifications. As I shall argue next, the prototypical specification in NLC should be the *set* stabilization problem. Although this problem does not capture all interesting control specifications (for instance, it does not encompass transient performance requirements), it is far richer than equilibrium stabilization, and it can be treated with minimal modifications of the concepts we already teach. Stability theory, for instance, can be taught seamlessly if one replaces equilibria by (controlled) invariant sets,[4] and indeed this idea emerges clearly in books such as [22], [23] and in many research papers devoted to stability theory.

We should not limit ourselves to replacing equilibria by invariant sets. An NLC course should show students how to convert practical specifications into set stabilization requirements. It should show the importance of set invariance in assessing the viability of a specification.

IV. Ideas for the Teaching of NLC Courses

In this Section I propose a few ideas for designing an NLC course, whether it be stability-based or geometry-based. The discussion concerns courses focusing on control design. As mentioned in Section II, geometric nonlinear control courses might spend a considerable amount of time on structural aspects of nonlinear control. The guidelines below may not be appropriate for such courses. The assumption in this discussion is that the course has a minimum duration of, say, 30 hours. Some of the ideas below might not be feasible for the one-week courses offered in some European summer schools. Finally, the below is a list of ideas, not a course outline.

Focus on a problem domain: Use a problem domain for motivation of course concepts. For instance, robotics, power systems, systems biology, spacecraft control. Possibly restrict the discussion to a particular class of systems, clarifying which course concepts can be generalized. For instance, if the problem area is robotics, one may restrict the focus to simple

[3]In the year 2000, the IEEE CDC conference had about 19 sessions dedicated to the topic of nonlinear systems and control. In the 2019 IEEE CDC (the last pre-pandemic edition of the conference), I counted only 7.

[4]Actually, compact sets. The stability and stabilization of non-compact sets presents additional challenges not present in the equilibrium version of these problems.

mechanical control systems [9]. Another, more general class of models with elegant connections to a variety of modelling domains is that of port-Hamiltonian systems [24], [25]. The instructor should not start their discussion from the differential equation $\dot{x} = f(x, u)$. Rather, some emphasis should be given on modelling within the chosen problem area.

Control specifications: Modelling goes hand in hand with control specifications. Place emphasis on typical control specifications from the chosen problem domain, formulating control problems and showing how to convert them into appropriate set stabilization problems. For instance, if the problem domain is distributed control, show how the synchronization of double-integrator robots corresponds to the stabilization of the diagonal subspace of their collective state space, while the stabilization of a rigid formation corresponds to the stabilization of a submanifold diffeomorphic to $\mathsf{SE}(3)$. Two motivating examples in the area of robotics are found in Sections V–VI.

Design guidelines: At the end of a presentation of a control methodology, present a procedural summary of the approach outlining the steps one has to follow to concretely solve the class of problems that the methodology addresses.

Problem-driven delivery: Use the problem domain as a means to introduce various course concepts. For instance, if the chosen problem domain were robotics, the zero dynamics manifold could be introduced using path following or virtual constraints, see Sections VI–VI.

Facilitate interdisciplinary interactions: If the course has a diverse graduate student population from engineering, physics, and mathematics, this is an asset that should be exploited. Post a list of ideas for final projects that students can use to formulate a project proposal. Make students present their proposals in mid-course, and their results at the end of the course. Consider posting team-based projects and facilitating competitions.

V. A Sample Lecture

This lecture introduces relative degree and zero dynamics via an example.

The Problem: Consider a bicycle travelling around a circular track of radius r depicted in Figure 1. We use Getz's bicycle model [26], in which the bicycle is a point-mass m moving in a vertical plane perpendicular to the direction of motion; see Figure 1. The bicycle has two control inputs: the steering torque and the force at p, the point of contact of the rear wheel with the ground. It is intuitively obvious that by appropriate choice of a steering feedback, the point p can be confined to lie on the circular track. Assuming this preliminary feedback has been applied, we are left with one control input, the force at the point p. We'll model the bicycle assuming p is constrained to lie on a circle of radius r.

Let b be the distance between p and the projection of the mass m along the inclined plane of the bicycle, and l be the distance between m and the ground along this plane, as in Figure 1. We assume that $r > l$, and ignore the moments of inertia of the wheels. With this simplification, the bicycle can be viewed as a pendulum rotating in a vertical plane orthogonal to the velocity $\dot{p}$. Its pivot point is offset from p at distance b along the direction tangent to the circle, where p moves along the circular track.

We place a right-handed inertial frame $oxyz$ at the centre of the circle, with z axis pointing upward. We let q_1 be the

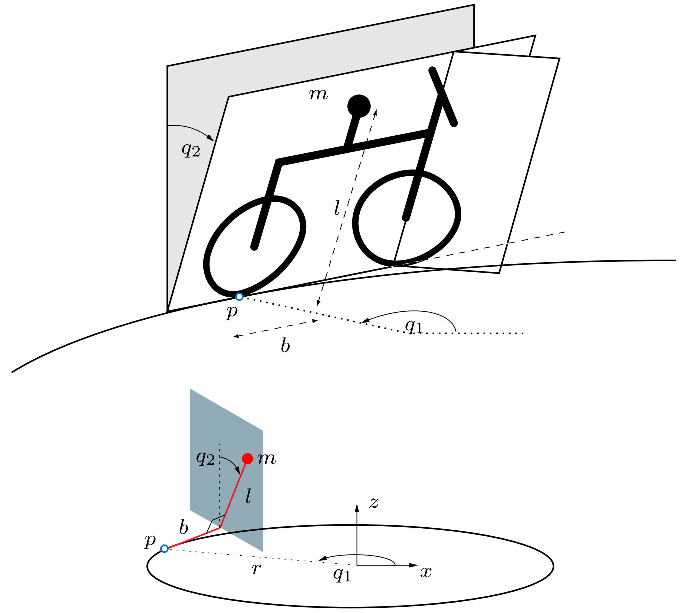

Fig. 1. Getz's bicycle on a circular track and its schematic representation as an offset pendulum.

angle of the point p measured from the x axis around the z axis according to the right-hand rule, and q_2 be the pendulum angle from the vertical line with the orientation displayed in the figure. Thus q_1 represents the position of the bicycle along the circle, while q_2 is its roll angle.

The generalized coordinates of the bicycle are $q = (q_1, q_2)$. The inertial coordinates of the mass m are

$$R(q) = \begin{bmatrix} r\cos(q_1) - b\sin(q_1) - l\cos(q_1)\sin(q_2) \\ r\sin(q_1) + b\cos(q_1) - l\sin(q_1)\sin(q_2) \\ l\cos(q_2) \end{bmatrix},$$

and the Lagrangian of the system is

$$L(q, \dot{q}) := \frac{1}{2}m\left\|\dot{R}(q, \dot{q})\right\|^2 - mgl\cos(q_2).$$

Assuming that the control force is tangent to the circle and given by $F = [-\sin(q_1)\ \ \cos(q_1)\ \ 0]^\top u/(r - l\sin(q_2))$, where $u \in \mathbb{R}$ is the control input, the corresponding generalized force is (see [27]) $(dR_q)^\top F = [u\ 0]^\top$, and the system model is

$$D(q)\ddot{q} + C(q, \dot{q})\dot{q} + \nabla_q P(q) = \begin{bmatrix} 1 \\ 0 \end{bmatrix} u, \qquad (1)$$

where, letting $s_{q_2} := \sin(q_2)$, $c_{q_2} := \cos(q_2)$, we have

$$D(q) = \begin{bmatrix} [c]m\left(l^2 s_{q_2}^2 + b^2 + r^2 - 2lrs_{q_2}\right) & blmc_{q_2} \\ blmc_{q_2} & ml^2 \end{bmatrix},$$

$$\nabla_q P(q) = -mlg\begin{bmatrix} 0 \\ s_{q_2} \end{bmatrix}.$$

The matrix $C(q, \dot{q})$ is obtained from $D(q)$ using the Christoffel symbols (see, e.g., [27]). The state of the bicycle is $x = (q, \dot{q})$ and the state space is $\mathscr{X} = ([\mathbb{R}]_{2\pi} \times [\mathbb{R}]_{2\pi}) \times (\mathbb{R} \times \mathbb{R})$, with $[\mathbb{R}]_{2\pi}$ denoting the real numbers mod 2π.

The control problem is to make the bicycle travel around the circle without falling over, and to achieve a constant steady-state speed independent of the initial conditions.

Notice that this control specification does not involve the stabilization of an equilibrium and does not lend itself to a linearization-based analysis. This problem is intrinsically nonlinear and demands nonlinear control tools.

The development that follows is inspired by the work in [28]. For the bicycle to travel at constant speed around the circle, it makes sense to require the roll angle, q_2, to be constant and contained in the interval $(0, \pi/2)$, so that the bicycle leans into the curve and the centripetal force balances the gravity force of the pendulum. Based on this observation, we enforce a *virtual constraint* imposing that the roll angle be constant and equal to some desired value $\theta \in (0, \pi/2)$.

Defining the output function[5] $h(x) = q_2 - \theta$, we want to design a controller sending the output $y(t) = h(x(t))$ to zero asymptotically. We also want to understand what happens to the closed-loop system in steady-state, when the roll angle is θ. Will the bicycle's speed converge to a steady-state value independent of initial conditions?

Controlled Invariance: The first step in our development is recognizing that there is no hope of asymptotically stabilizing the zero level set of the output, the set $h^{-1}(0) \subset \mathscr{X}$. Indeed, a necessary condition for a set to be stable is that the set be invariant, and the set $h^{-1}(0)$ cannot be rendered invariant via feedback (in other words, $h^{-1}(0)$ is not controlled invariant).

To understand why this is the case, suppose the initial configuration of the bicycle is $q(0) = (0, \theta) \in h^{-1}(0)$, and that $\dot{q}(0) = (0, 1)$. Then, $\dot{y}(0) = \dot{q}_2 = 1$, implying that $y(t) > 0$ for small $t > 0$ so the roll angle of the bicycle will increase past θ after a short amount of time, no matter what control input u is applied.

Thus the set $h^{-1}(0)$ is not controlled invariant and as such, it cannot be stabilized. One could make the set $h^{-1}(0)$ attractive (but unstable), but this would be undesirable in practice, as attractivity is a fragile property that is destroyed by arbitrarily small model uncertainties.

The question then is, what is the *largest* controlled invariant subset of $h^{-1}(0)$? To answer this question, we perform a thought experiment. Suppose $x(t)$ is a solution confined within the set $h^{-1}(0)$, i.e., such that $y(t) = h(x(t)) \equiv 0$. Then, necessarily, $\dot{y}(t) \equiv 0$, or $\dot{q}_2(t) \equiv 0$. In order for $x(t)$ to remain in the set $h^{-1}(0)$, $x(t)$ must be contained in the subset of $h^{-1}(0)$

$$\begin{aligned}\mathcal{Z} &:= \left\{x \in \mathscr{X} : h(x) = L_f h(x) = 0\right\} \\ &= \{x \in \mathscr{X} : q_2 = \dot{q}_2 = 0\}.\end{aligned}$$

Is the set $\mathcal{Z}$ controlled invariant? We show that the answer is *yes*. We have

$$\begin{aligned}y &= q_2 - \theta \\ \dot{y} &= \dot{q}_2 \\ \ddot{y} &= \ddot{q}_2 = \alpha(x) + \beta(x)u,\end{aligned}$$

where $\alpha(x)$ is a suitable smooth function and

$$\beta(x) = \frac{-b c_{q_2}}{ml\Big(b^2 c_{q_2}^2 + \big(r - l s_{q_2}\big)^2\Big)}.$$

[5]Recall that the configuration variable q_2 lives in $[\mathbb{R}]_{2\pi}$. The chosen output h then is not real-valued. We will ignore this fact to keep the presentation simple, but a better choice of output would be $h(x) = \sin(q_2 - \theta)$. With this choice, $h(x) = 0$ if $q_2 = [\theta]_{2\pi}$ or $[\theta + \pi]_{2\pi}$. This apparent inconvenience does not actually affect the results that follow.

Since $r > l$, the denominator of β is positive, and $\beta(x) \neq 0$ on the set $\mathcal{U} := \{x \in \mathscr{X} : \cos(q_2) \neq 0\}$. Note that $\mathcal{Z} \subset \mathcal{U}$.

Thus the control input u appears in the second time derivative of the output along solutions of the control system, and the coefficient of u is nonzero on the set $\mathcal{Z}$. We say that the cart-pendulum system with output $y = h(x)$ has a *well-defined relative degree 2*, the integer 2 signifying the order of differentiation of y before u appears nonsingularly [7].

The feedback

$$u^\star(x) = -\frac{\alpha(x)}{\beta(x)} \tag{2}$$

is well-defined on the set $\mathcal{Z}$ and it gives $\ddot{y} \equiv 0$ on this set. As we show in a moment, this means that $\mathcal{Z}$ is rendered invariant by $u^\star$. Indeed, if $x(0) \in \mathcal{Z}$, then $y(0) = \dot{y}(0) = 0$. Since $\ddot{y}(t) = 0$, $\dot{y}(t)$ is constant, and therefore $\dot{y}(t) \equiv \dot{y}(0) = 0$. In turn, $\dot{y}(t) \equiv 0$ implies that $y(t)$ is constant, or $y(t) \equiv y(0) = 0$. In conclusion, if $x(0) \in \mathcal{Z}$, then $x(t) \in \mathcal{Z}$ for all $t \in \mathbb{R}$ so the feedback $u^\star(x)$ in (2) renders $\mathcal{Z}$ invariant, and $\mathcal{Z}$ is controlled invariant. We call $u^\star(x)$ the *friend* of $\mathcal{Z}$, in the terminology of Wonham [29].

The construction above makes it clear that $\mathcal{Z}$ is the maximal controlled invariant subset of $h^{-1}(0)$. We call $\mathcal{Z}$ the *zero dynamics manifold* [20], [21] of the bicycle with output $y = h(x)$.

The physical significance of $\mathcal{Z}$ is this. Since there is no hope of asymptotically stabilizing the set $h^{-1}(0)$, the next best thing is to asymptotically stabilize its maximal controlled invariant subset, and we've shown that this subset is $\mathcal{Z}$. Our control specification then is the asymptotic stabilization of the zero dynamics manifold $\mathcal{Z}$.

Asymptotic Stabilization of the Zero Dynamics Manifold: We return to the identity found earlier,

$$\ddot{y} = \alpha(x) + \beta(x)u,$$

where $\beta(x) \neq 0$ on the set $\mathcal{U}$. The feedback transformation on $\mathcal{U}$ given by

$$u(x) = \frac{1}{\beta(x)}(-\alpha(x) + w),$$

where $w \in \mathbb{R}$ is a new control input, gives $\ddot{y} = w$, an LTI controllable system with input w (a double-integrator). Setting $w = -K_p y - K_d \dot{y}$, with $K_p, K_d > 0$, stabilizes the origin $(y, \dot{y}) = (0, 0)$ of the LTI system. The resulting controller in x coordinates,

$$\bar{u}(x) = \frac{1}{\beta(x)}\big(-\alpha(x) - K_p h(x) - K_d L_f h(x)\big), \tag{3}$$

where $h(x) = q_2 - \theta$ and $L_f h(x) = \dot{q}_2$, is well-defined on $\mathcal{U}$. The feedback $\bar{u}(x)$ is called an *input-output feedback linearizing controller*, with reference to the fact that it renders the dynamics from the input w to the output y LTI.

Returning to x coordinates, recall that $\mathcal{Z} = \{x : h(x) = L_f h(x) = 0\}$. If the closed-loop system given by bicycle with feedback controller $\bar{u}(x)$ has no finite escape times,[6] the asymptotic stability of the equilibrium $(y, \dot{y}) = (0, 0)$ for the double-integrator implies the local asymptotic stability of the set $\mathcal{Z}$, solving our control problem.

[6]The issue of finite escape times can be addressed by investigating the closed-loop dynamics using the normal form for input-output feedback linearization, a topic the instructor can make reference to if they plan to teach it.

Zero Dynamics: Since, on $\mathcal{Z}$, $h(x) = L_f h(x) = 0$, the feedback $\bar{u}(x)$ in (3) reduces to the friend $u^\star(x)$ in (2) rendering $\mathcal{Z}$ invariant, and so $\mathcal{Z}$ is an invariant set for the closed-loop system. The dynamics on $\mathcal{Z}$ of the bicycle with feedback $u^\star(x)$ are called *the zero dynamics* [21].

The zero dynamics characterize the steady-state behaviour of all bounded solutions $x(t)$ such that $h(x(t)) \to 0$. To see why this is the case, we first recall that if $x(t)$ is a bounded solution of the closed-loop system, then by the Birkhoff theorem (see [4, Lemma 4.1]), the positive limit set $L^+(x(0))$ of $x(t)$ is nonempty, compact, and invariant. Moreover, $x(t)$ converges asymptotically to $L^+(x(0))$. Using standard arguments (e.g., the ones used in the proof of the Krasovskii-LaSalle invariance principle) and the fact that $h(x(t)) \to 0$, one can show that $L^+(x(0)) \subset h^{-1}(0)$. Since $L^+(x(0))$ is an invariant set for the closed-loop system contained in $h^{-1}(0)$, and since the zero dynamics manifold $\mathcal{Z}$ is the maximal controlled invariant subset of $h^{-1}(0)$, $L^+(x(0)) \subseteq \mathcal{Z}$. Any bounded solution $x(t)$ of the closed-loop system converges to $L^+(x(0)) \subseteq \mathcal{Z}$. Thus, the closed-loop dynamics on $\mathcal{Z}$ characterize the steady-state behaviour of any bounded solution $x(t)$ such that $h(x(t)) \to 0$, as claimed.

To find the dynamics on $\mathcal{Z}$, we set $u = \bar{u}(x)$ in (1) and impose $h(x) = L_f h(x) = 0$, or $q_2 = \theta, \dot{q}_2 = 0$. Substituting these identities in the $\ddot{q}_1$ equation, we arrive at

$$\ddot{q}_1 = \frac{g\tan(\theta)}{b} - \frac{r - l\sin(\theta)}{b}\dot{q}_1^2. \tag{4}$$

This second-order differential equation represents the zero dynamics of the bicycle, i.e., the dynamics on $\mathcal{Z}$, in the following sense: any solution $x(t)$ entirely contained in $\mathcal{Z}$ has the form $x(t) = (q_1(t), \dot{q}_1(t), \theta, 0)$, where $(q_1(t), \dot{q}_1(t))$ is a solution of (4). Letting $\omega := \dot{q}_1$, we rewrite the above differential equation as

$$\dot{\omega} = \frac{g\tan(\theta)}{b} - \frac{r - l\sin(\theta)}{b}\omega^2.$$

Letting

$$\bar{\omega} := \left(\frac{g\tan(\theta)}{r - l\sin(\theta)}\right)^{1/2},$$

there are two equilibria at $\omega = \pm\bar{\omega}$. They correspond to closed orbits $\mathcal{O}^\pm := \{x \in \mathcal{Z} : \dot{q}_1 = \pm\bar{\omega}\}$ of the zero dynamics (4). On $\mathcal{O}^+$, the point p of the bicycle moves forward along the circle with speed $\|\dot{p}\| = \bar{\omega}$, while on $\mathcal{O}^-$ it moves backward with the same speed.

If $\omega^2 > \bar{\omega}^2$ then $\dot{\omega} < 0$, while if $\omega^2 < \bar{\omega}^2$ we have $\dot{\omega} > 0$. This implies that for all $\omega(0) > -\bar{\omega}$, the solution $\omega(t)$ converges to $\bar{\omega}$, while for $\omega(0) < -\bar{\omega}$, $\omega(t) \to -\infty$ in finite time. Thus the closed orbit $\mathcal{O}^+$ is a stable limit cycle for the zero dynamics, while the closed orbit $\mathcal{O}^-$ is an unstable limit cycle. The domain of attraction of $\mathcal{O}^+$ is $\{x \in \mathcal{Z} : \dot{q}_1 > -\bar{\omega}\}$. A phase portrait of the orbits of (4) with parameters $(r, b, l, \theta) = (100, 1, 1, \pi/6)$ is found in Figure 2.

The physical meaning of these considerations is as follows. When initialized on the zero dynamics manifold with speed $\dot{q}_1 > -\bar{\omega}$, the bicycle will traverse the circular track and reach a steady-state speed $\bar{\omega}$ independent of initial conditions. By changing the roll angle θ we change the steady-state speed $\bar{\omega}$. Since the function $(0, \pi/2) \to \mathbb{R}_{>0}, \theta \mapsto \bar{\omega}$ is surjective, we can arbitrarily assign the steady-state speed via $\theta \in (0, \pi/2)$. When initialized off $\mathcal{Z}$ but sufficiently close to it, the bicycle will converge to it and display one of two behaviours. It will either have a finite escape time if it approaches the part of $\mathcal{Z}$ not included in the domain of attraction of $\mathcal{O}^+$, or it will converge to $\mathcal{O}^+$, therefore meeting our control specification. One can show that $\mathcal{O}^+$ is actually locally asymptotically stable for the closed-loop system.

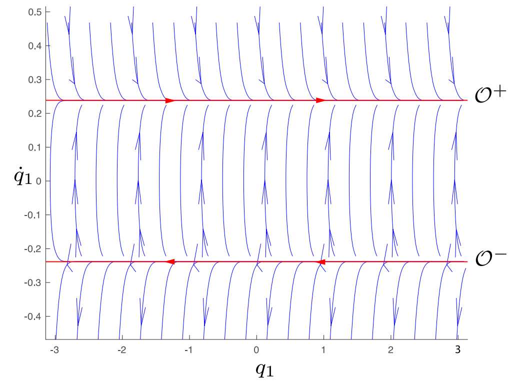

Fig. 2. The zero dynamics of the bicycle. The vertical lines $q_1 = \pm\pi$ are identified. The two red lines represent closed orbits, a stable and an unstable limit cycle.

At this point the instructor ends the lecture with a point-format summary of concepts. The next lecture will be devoted to formalizing the ideas that emerge in the example above. The level of rigour and the extent of the formalization will depend on the style the instructor chooses to adopt.

Summary of Concepts:

- A control specification was encoded into an output zeroing problem.
- The control system possesses a well-defined relative degree r, i.e., the control input appears nonsingularly after r time differentiations of the output along solutions.
- Relative degree allowed us to compute the maximal controlled invariant subset of the zero level set of the output. This is the zero dynamics manifold $\mathcal{Z}$. To meet the control specification, we wanted to asymptotically stabilize $\mathcal{Z}$.
- We asymptotically stabilized $\mathcal{Z}$ by means an input-output feedback linearizing controller.
- $\mathcal{Z}$ is an invariant set for the closed-loop system, and solutions on this set characterize the steady-state behaviour of all bounded solutions whose output signals tend to zero.
- The zero dynamics is the vector field representing the dynamics on $\mathcal{Z}$ obtained using the friend $u^\star$ rendering $\mathcal{Z}$ invariant.

VI. A Course Assignment

This assignment builds upon the previous lecture. Consider the kinematic unicycle model

$$\begin{aligned} \dot{x}_1 &= v\cos(x_3) \\ \dot{x}_2 &= v\sin(x_3) \\ \dot{x}_3 &= u, \end{aligned} \tag{5}$$

where $v > 0$ is a fixed parameter representing the unicycle speed and $u \in \mathbb{R}$ is the control input. The states $(x_1, x_2) \in \mathbb{R} \times \mathbb{R}$ represents the position of the unicycle in the plane, and $x_3 \in [\mathbb{R}]_{2\pi}$ is the heading angle. We denote by $x = (x_1, x_2, x_3)$ the state, and by $\mathscr{X}$ the state space.

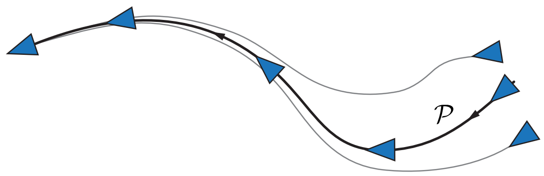

Fig. 3. The path following problem.

Let $\mathcal{P}$ be a curve in $\mathbb{R} \times \mathbb{R}$, represented implicitly as $\mathcal{P} = \{(x_1, x_2) \in \mathbb{R} \times \mathbb{R} : \gamma(x_1, x_2) = 0\}$, where $\gamma : \mathbb{R} \times \mathbb{R} \to \mathbb{R}$ is C^1 function whose Jacobian matrix is nonvanishing on $\mathcal{P}$, and therefore by continuity of the Jacobian, nonvanishing in a neighbourhood $\mathcal{U}$ of $\mathcal{P}$. Given a real number $c \in \mathbb{R}$, we define the level set of γ, $\mathcal{P}_c := \{(x_1, x_2) \in \mathbb{R} \times \mathbb{R} : \gamma(x_1, x_2) = c\}$. With this notation, $\mathcal{P}_0 = \mathcal{P}$.

We define a few useful geometric vectors. Consider a point $(x_1, x_2) \in \mathcal{U}$ and let $c := \gamma(x_1, x_2)$.

- The unit normal vector to $\mathcal{P}_c$ at (x_1, x_2) is $\mathbf{n}(x_1, x_2) := \nabla_{(x_1,x_2)}\gamma / \|\nabla_{(x_1,x_2)}\gamma\|$.
- The unit tangent vector to $\mathcal{P}_c$ at (x_1, x_2) is $\mathbf{t}(x_1, x_2)$ given by the *clockwise* rotation of $n(x_1, x_2)$ by angle $\pi/2$.
- The heading vector of the unicycle is $\mathbf{h}(x_3) = [\cos(x_3)\ \ \sin(x_3)]^\top$.

The tangent vector $\mathbf{t}$ defined above endows the curve $\mathcal{P}$ with an orientation. One could change the orientation of $\mathcal{P}$ by changing the sign of $\mathbf{t}$ (i.e., rotating $\mathbf{n}$ *counterclockwise*).

Path following problem (PFP): find a smooth feedback controller $\bar{u}(x)$ for (5) meeting the following specifications (see Figure 3):

(i) *(Path invariance)* When the unicycle's position is initialized on $\mathcal{P}$ with heading tangent to $\mathcal{P}$ and consistent with the orientation of $\mathcal{P}$, i.e., $\mathbf{h}(x_3) = \mathbf{t}(x_1, x_2)$, the unicycle remains on $\mathcal{P}$ for all $t \in \mathbb{R}$.

(ii) *(Convergence)* For suitable initial conditions near the path $\mathcal{P}$ with heading not too far from the tangent heading consistent with the orientation, the unicycle asymptotically converges to $\mathcal{P}$, traversing it with nonzero speed in the direction or the orientation of $\mathcal{P}$.

(iii) *(Stability)* If the unicycle is initialized near $\mathcal{P}$ with heading vector close to $\mathbf{t}$, then the unicycle remains close to $\mathcal{P}$, and its heading remains close to $\mathbf{t}$.

Question 1: define the subset of the state space where the unicycle is on $\mathcal{P}$ with heading tangent to $\mathcal{P}$ and consistent with the orientation,

$$\mathcal{Z}^+ := \{x \in \mathcal{X} : \gamma(x_1, x_2) = 0, \mathbf{h}(x_3) = \mathbf{t}(x_1, x_2)\}. \tag{6}$$

Prove that PFP is equivalent to requiring that $\mathcal{Z}^+$ be asymptotically stable.

Solution outline: Requirement (i) of the PFP means that the controller $\bar{u}$ must render $\mathcal{Z}^+$ an invariant set, while requirement (ii) means that the controller must render $\mathcal{Z}^+$ locally attractive. Finally, requirement (iii) means that if the unicycle is initialized near $\mathcal{Z}^+$, it must stay near $\mathcal{Z}^+$, i.e., $\mathcal{Z}^+$ must be a stable set for the closed-loop system. Since set stability implies set invariance, the PFP is mathematically equivalent to the local asymptotic stabilization of the set $\mathcal{Z}^+$.

Question 2: Show that the unicycle model (5) with input u and output function $h(x) = \gamma(x_1, x_2)$ has a well-defined relative degree 2 everywhere on $\mathcal{X}$ except for a "singularity" set $\mathcal{S} \subseteq \mathcal{X}$. Explain what is the physical configuration of the unicycle when its state is in $\mathcal{S}$.

Solution outline: Since $\ddot{y} = a(x) + \langle \nabla_{(x_1,x_2)}\gamma, \mathbf{h}(x_3)\rangle v\, u$, the system has relative degree 2 on the set $\mathcal{X} \setminus \mathcal{S}$, with $\mathcal{S} = \{x \in \mathcal{X} : \langle \nabla_{(x_1,x_2)}\gamma, \mathbf{h}(x_3)\rangle = 0\}$. On $\mathcal{S}$, either $\nabla_{(x_1,x_2)}\gamma = 0$, or the unicycle heading is parallel to the normal vector $\mathbf{n}$, i.e., the unicycle points perpendicularly towards or perpendicularly away from the curve $\mathcal{P}$.

Question 3: Show that the zero dynamics manifold, $\mathcal{Z}$, of the unicycle model (5) with output function $h(x)$ defined above is the union of two disjoint sets, one of which is the set $\mathcal{Z}^+$ in (6). Explain what is the physical configuration of the unicycle when its state is on the other set. Show that $\mathcal{Z} \cap \mathcal{S} = \emptyset$.

Solution outline: the system has relative degree 2 on $\mathcal{X} \setminus \mathcal{S}$, and $\mathcal{S} \cap \{x : h(x) = L_f h(x) = 0\} = \emptyset$, thus the zero dynamics manifold is

$$\begin{aligned}\mathcal{Z} &= \left\{x \in \mathcal{X} : h(x) = L_f h(x) = 0\right\} \\ &= \{x \in \mathcal{X} : \gamma(x_1, x_2) = 0, \langle \mathbf{n}(x_1, x_2), \mathbf{h}(x_3)\rangle = 0\} \\ &= \mathcal{Z}^+ \cup \mathcal{Z}^-,\end{aligned}$$

where $\mathcal{Z}^- = \{x \in \mathcal{X} : \gamma(x_1, x_2) = 0, \mathbf{h}(x_3) = -\mathbf{t}(x_1, x_2)\}$. It's easy to show that $\mathcal{Z}^+ \cap \mathcal{Z}^- = \emptyset$. On $\mathcal{Z}^-$, the unicycle is on $\mathcal{P}$, with heading tangent to $\mathcal{P}$ but pointing in the direction opposite the orientation of $\mathcal{P}$.

Now the instructor can ask students to design an input-output linearizing controller and argue with varying levels of rigour that the controller asymptotically stabilizes $\mathcal{Z}$, $\mathcal{Z}^+$, and $\mathcal{Z}^-$. In particular, the controller solves the PFP. Or else, ask them to carry out extensive numerical simulations and collect statistics about which initial conditions lead to failure (i.e., the state trajectory hits $\mathcal{S}$), converge to $\mathcal{Z}^+$, or converge to $\mathcal{Z}^-$.

VII. Conclusion

This letter discusses graduate education in NLC and suggests replacing the current focus on equilibrium stability and stabilization by a new focus on *set* stability and stabilization. Together with this change of focus, this letter proposes placing emphasis on control specifications inspired by specific problem domains. When teaching zero dynamics, for instance, one could introduce ideas using examples from robotics, and this letter proposes two such examples.

Some of the challenges I highlighted in this letter are not relegated to educational aspects of NLC. There would be much to say about the current state of the discipline and of control theory at large. It would be worthwhile for the research community to assess where the discipline is heading, and for that we need to create new avenues of discussion.

Finally, the proposals advanced in this letter for teaching NLC courses assume an emphasis on control design and do not pertain to those courses placing emphasis on structural aspects of NLC. Those courses would merit a separate discussion.

References

[1] A. Loría and E. Panteley, "Stability, told by its developers," *Advanced Topics in Control Theory* (Lecture Notes in Control and Information Sciences), A. Loría, F. Lamnabhi-Lagarrigue, E. Panteley, Eds. London, U.K.: Springer, 2006.

[2] J.-L. Lagrange, *Mécanique Analytique*, 1st ed., Chez la veuve Desaint, 1788.

[3] A. Lyapunov, "The general problem of motion stability," in *Kharkovskoye Matematicheskoe Obshchestvo*, vol. 11, 1892.

[4] H. K. Khalil, *Nonlinear Systems*, 3rd ed. Upper Saddle River, NJ, USA: Prentice Hall, 2002.

[5] J.-J. Slotine and W. Li, *Applied Nonlinear Control*, vol. 199. Englewood Cliffs, NJ, USA: Prentice Hall, 1991.

[6] M. Vidyasagar, *Nonlinear Systems Analysis*. Philadelphia, PA, USA: SIAM, 2002.
[7] A. Isidori, *Nonlinear Control Systems*, 3rd ed. London, U.K.: Springer-Verlag, 1995.
[8] H. Nijmeijer and A. V. der Schaft, *Nonlinear Dynamical Control Systems*, vol. 175. New York, NY, USA: Springer, 1990.
[9] F. Bullo and A. D. Lewis, *Geometric Control of Mechanical Systems* (Texts in Applied Mathematics), vol. 49. New York, NY, USA: Springer, 2005.
[10] V. Jurdjevic, *Geometric Control Theory*. Cambridge, U.K.: Cambridge Univ. Press, 1997.
[11] A. Agrachev and Y. Sachkov, *Control Theory From the Geometric Viewpoint*, vol. 2. Berlin, Germany: Springer, 2004.
[12] A. Roza, M. Maggiore, and L. Scardovi, *Distributed Coordination Theory for Robot Teams*. Cham, Switzerland: Springer, 2022. [Online]. Available: https://dx.doi.org/10.1007/978-3-030-96087-2
[13] E. Westervelt, J. W. Grizzle, and D. E. Koditschek, "Hybrid zero dynamics of planar biped walkers," *IEEE Trans. Autom. Control*, vol. 48, no. 1, pp. 42–56, Jan. 2003.
[14] R. Paul, "Modelling, trajectory calculation and servoing of a computer controlled arm," Ph.D. dissertation, Dept. Comput. Sci., Stanford Univ., Stanford, CA, USA, 1972.
[15] B. Markiewicz, "Analysis of the computed torque drive method and comparison with conventional position servo for a computer-controlled manipulator," Jet Propulsion Lab., Pasadena, CA, Rep. NASA TM 33-601, 1973.
[16] R. Brockett, "Feedback invariants for nonlinear systems," *IFAC Proc. Vol.*, vol. 11, no. 1, pp. 1115–1120, 1978.
[17] R. Su, "On the linear equivalents of nonlinear systems," *Syst. Control Lett.*, vol. 2, no. 1, pp. 48–52, 1982.
[18] B. Jakubczyk and W. Respondek, "On Linearization of control systems," *Bull. Acad. Polon. Sci. Sci. Math*, vol. 28, nos. 9–10, pp. 517–522, 1980.
[19] L. Hunt, R. Su, and G. Meyer, "Design for multi-input nonlinear systems," in *Differential Geometric Control Theory*, R. Brockett, R. Millman, and H. Sussmann, Eds. Boston, MA, USA: Birkhäuser, 1983, pp. 268–298.
[20] C. I. Byrnes and A. Isidori, "A frequency domain philosophy for nonlinear systems, with applications to stabilization and to adaptive control," in *Proc. 23rd IEEE Conf. Decis. Control*, 1984, pp. 1569–1573.
[21] C. I. Byrnes and A. Isidori, "Asymptotic stabilization of minimum phase nonlinear systems," *IEEE Trans. Autom. Control*, vol. 36, no. 10, pp. 1122–1137, Oct. 1991.
[22] N. P. Bathia and G. P. Szegö, *Stability Theory of Dynamical Systems*. Berlin, Germany: Springer-Verlag, 1970.
[23] R. Goebel, R. Sanfelice, and A. Teel, *Hybrid Dynamical Systems: Modeling, Stability, and Robustness*. Princeton, NJ, USA: Princeton Univ. Press, 2012.
[24] R. Ortega, A. V. D. Schaft, B. Maschke, and G. Escobar, "Interconnection and damping assignment passivity-based control of port-controlled Hamiltonian systems," *Automatica*, vol. 38, no. 4, pp. 585–596, 2002.
[25] A. Van der Schaft, *L_2-Gain and Passivity Techniques in Nonlinear Control*. (Communications and Control Engineering), 3rd ed. Cham, Switzerland: Springer, 2017.
[26] N. Getz, "Control of balance for a nonlinear nonholonomic non- minimum phase model of a bicycle," in *Proc. Amer. Control Conf.*, 1994, pp. 148–151.
[27] M. W. Spong, S. Hutchinson, and M. Vidyasagar, *Robot Modeling and Control*. Hoboken, NJ, USA: Wiley, 2020.
[28] L. Consolini and M. Maggiore, "Control of a bicycle using virtual holonomic constraints," *Automatica*, vol. 49, pp. 2381–2839, Sep. 2013.
[29] W. Wonham, *Linear Multivariable Control: A Geometric Approach*, 3rd ed. New York, NY, USA: Springer-Verlag, 1985.