ECE 517:
Nonlinear and Adaptive Control

Fall '09 Lecture Notes

Daniel Liberzon

November 29, 2009

Disclaimers

Don't print future lectures in advance as the material is always in the process of being updated.

These lecture notes are posted for class use only.

This is a very rough draft which contains many errors.

I don't always give proper references to sources from which results are taken. A lack of reference does not mean that the result is original. In fact, all results presented in these notes (with possible exception of some simple examples) were borrowed from the literature and are not mine.

Contents

Lecture 1 (Mon Aug 31)

1 Introduction

The meaning of "nonlinear" should be clear, even if you only studied linear systems so far (by exclusion).

The meaning of "adaptive" is less clear and takes longer to explain.

From www.webster.com:

Adaptive: showing or having a capacity for or tendency toward *adaptation*.

Adaptation: the act or process of *adapting*.

Adapt: to become *adapted*.

Perhaps it's easier to first explain the class of problems it studies: modeling uncertainty. This includes (but is not limited to) the presence of unknown parameters in the model of the plant.

There are many specialized techniques in adaptive control, and details of analysis and design tend to be challenging. We'll try to extract fundamental concepts and ideas, of interest not only in adaptive control. The presentation of adaptive control results will mostly be at the level of examples, not general theory.

The pattern will be: general concept in nonlinear systems/control, followed by its application in adaptive control. Or, even better: a motivating example/problem in adaptive control, then the general treatment of the concept or technique, then back to its adaptive application. Overall, the course is designed to provide an *introduction* to further studies both in nonlinear systems and control and in adaptive control.

1.1 Motivating example

Example 1 Consider the scalar system

$$\dot{x} = \theta x + u$$

where x is state, u is control, and θ is an unknown fixed parameter.

A word on notation: There's no consistent notation in adaptive control literature for the true value of the unknown parameters. When there is only one parameter, θ is a fairly standard symbol. Sometimes it's denoted as θ^* (to further emphasize that it is the actual value of θ). In other sources, p^* is used. When there are several unknown parameters, they are either combined into a vector (θ, θ^*, p^*, etc.) or written individually using different letters such as a, b, and so on. Estimates of the unknown parameters commonly have hats over them, e.g., $\hat{\theta}$, and estimation errors commonly have tildas over them, e.g., $\tilde{\theta} = \hat{\theta} - \theta$.

Goal: regulation, i.e., make $x(t) \to 0$ as $t \to \infty$.

If $\theta < 0$, then $u \equiv 0$ works.

If $\theta > 0$ but is known, then the feedback law

$$u = -(\theta + 1)x$$

gives $\dot{x} = -x \Rightarrow x \to 0$. (Instead of $+1$ can use any other positive number.)

But if (as is the case of interest) θ is unknown, this u is not implementable.

$\longrightarrow$ Even in this simplest possible example, it's not obvious what to do.

Adaptive control law:

$$\dot{\hat{\theta}} = x^2 \tag{1}$$

$$u = -(\hat{\theta} + 1)x \tag{2}$$

Here (1) is the *tuning law*, it "tunes" the feedback gain.

Closed-loop system:

$$\begin{aligned} \dot{x} &= (\theta - \hat{\theta} - 1)x \\ \dot{\hat{\theta}} &= x^2 \end{aligned}$$

Intuition: the growth of $\hat{\theta}$ dominates the linear growth of x, and eventually the feedback gain $\hat{\theta} + 1$ becomes large enough to overcome the uncertainty and stabilize the system.

Analysis: let's try to find a Lyapunov function.

If we take

$$V := \frac{x^2}{2}$$

then its derivative along the closed-loop system is

$$\dot{V} = (\theta - \hat{\theta} - 1)x^2$$

and this is not guaranteed to be negative.

Besides, V should be a function of both states of the closed-loop system, x and $\hat{\theta}$.

Actually, with the above V we can still prove stability, although analysis is more intricate. We'll see this later.

Let's take

$$V(x, \hat{\theta}) := \frac{x^2}{2} + \frac{(\hat{\theta} - \theta)^2}{2} \tag{3}$$

The choice of the second term reflects the fact that in principle, we want to have an asymptotically stable equilibrium at $x = 0$, $\hat{\theta} = \theta$. In other words, we can think of $\hat{\theta}$ as an *estimate* of θ. However, the control objective doesn't explicitly require that $\hat{\theta} \to \theta$.

With V given by (3), we get

$$\dot{V} = (\theta - \hat{\theta} - 1)x^2 + (\hat{\theta} - \theta)x^2 = -x^2 \tag{4}$$

Is this enough to prove that $x(t) \to 0$?

Recall:

Theorem 1 (Lyapunov) *Let V be a positive definite $\mathcal{C}^1$ function. If its derivative along solutions satisfies*

$$\dot{V} \le 0 \tag{5}$$

everywhere, then the system is stable. If

$$\dot{V} < 0 \tag{6}$$

everywhere (except at the equilibrium being studied), then the system is asymptotically stable. If in the latter case V is also radially unbounded (i.e., $V \to \infty$ as the state approaches ∞ along any direction), then the system is globally asymptotically stable.

From (4) we certainly have (5), hence we have stability (in the sense of Lyapunov). In particular, both x and $\hat{\theta}$ remain bounded for all time (by a constant depending on initial conditions).

On the other hand, we don't have (6) because $\dot{V} = 0$ for $(x, \hat{\theta})$ with $x = 0$ and $\hat{\theta}$ arbitrary. It seems plausible that at least convergence of x to 0 should follow from (4). This is indeed true, but proving this requires knowing a precise result about *weak* (nonstrictly decreasing) Lyapunov functions. We will learn/review such results and will then finish the example. □

Some observations from the above example:

- Even though the plant is linear, the control is *nonlinear* (because of the square terms). To analyze the closed-loop system, we need nonlinear analysis tools.
- The control law is *dynamic* as it incorporates the tuning equation for $\hat{\theta}$. Intuitively, this equation "learns" the unknown value of θ, providing estimates of θ.
- The standard Lyapunov stability theorem is not enough, and we need to work with a weak Lyapunov function. As we will see, this is typical in adaptive control, because the estimates ($\hat{\theta}$) might not converge to the actual parameter values (θ). We will discuss techniques for making the parameter estimates converge. However, even without this we can achieve regulation of the state x to 0 (i.e., have convergence of those variables that we care about).

So, what is adaptive control?

We can think of the tuning law in the above example as an *adaptation* block in the overall system—see figure. (The diagram is a bit more general than what we had in the example.)

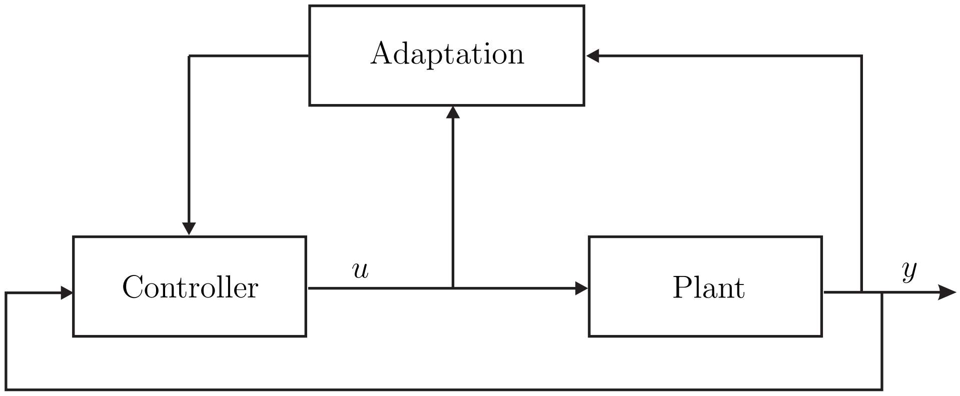

A priori, this is no different from any other (non-adaptive) dynamic feedback. But the main feature of the adaptive controller is that it achieves the control objective (regulation) despite *large-scale uncertainty* associated with θ (which can be any real number). This is to be contrasted with *robust control* where uncertainty range is usually bounded and small.

So, one may call a controller *adaptive* if it can handle such systems, i.e., if it can "adapt" to large uncertainties. Ideally we want an adaptive controller to handle more than just constant unknown parameters: parameters may vary with time, there may be unmodeled dynamics, noise and disturbance inputs, etc.

Certainty equivalence principle:

The way we designed the controller in Example 1 was:

- If θ were known, we would use $u = -(\theta + 1)x$.
- Since θ was unknown, we introduced its estimate, $\hat{\theta}$.
- Then we used the control $u = -(\hat{\theta} + 1)x$. I.e., we used the estimate for control purposes, pretending that it is correct, even though it may not be (and it works, at least for this example).

Adaptive controllers designed using the above logic are called certainty equivalence controllers. This means essentially that estimator design and controller design were decoupled. A vast majority of adaptive control techniques are based on certainty equivalence principle, although in some situations it may be desirable to rethink this and to design the controller by explicitly taking parameter estimation into account. (More on this later.)

1.2 Course logistics

⟶ On those Wednesdays when there is a Decision and Control seminar, we will start about 5 minutes late.

Prerequisite is ECE 515 (Linear Systems). Some knowledge of nonlinear systems (such as Lyapunov stability, actually covered to some extent in ECE 515) is a plus. If everything in the above discussion was familiar to you, you should be OK.

There is no single textbook for this course. Lecture notes will be posted on the class website. Material for the lecture notes is drawn from several sources.

- Adaptive control of linear plants:

  Ioannou and Sun, *Robust Adaptive Control*, Prentice-Hall, 1996. Out of print, available on-line (see class website).

  Ioannou and Fidan, *Adaptive Control Tutorial*, SIAM, 2006. This is an updated (and somewhat simplified) version of the first book.

- Nonlinear systems:

  Khalil, *Nonlinear Systems*, Prentice-Hall, 2002 (third edition). This is a standard and very good text, which is also the main text for ECE 528 (Nonlinear Systems). It has some adaptive control examples too (see index). If you're planning to take ECE 528, get it now.

- Adaptive control of nonlinear plants:

  Krstic, Kanellakopoulos and Kokotović, *Nonlinear and Adaptive Control Design*, Wiley, 1995. (Referred to as "KKK book" below.) Contains some more advanced adaptive control material, and covers some nonlinear systems and control theory concepts as well.

- Some material is also drawn from lecture notes on adaptive control by A. S. Morse (communicated privately to the instructor). These are in turn based on research articles, I can provide the references upon request.

For outline of the topics, see the table of contents. On average, 1–1.5 weeks will be spent on each topic (except parameter estimation which is longer).

⟶ Since everything discussed in class will be in the lecture notes posted on the web (well, except perhaps some pictures I might draw on the board to explain something better), you don't have to worry about copying things down. Instead, try to understand and participate in the discussion as much as possible.

Grading scheme: Homework—50%, Final Project—50%.

Homework: There will be about 4 problem sets. Some MATLAB simulations, to validate the designs discussed in class and study their robustness properties. Theoretical questions as well.

Project: Find another concept useful in adaptive control. Present the concept and its adaptive application. At least one of the two should be interesting/nontrivial. (These guidelines are somewhat flexible.) Prepare a self-contained presentation (.ppt, .pdf) understandable to other students. Length of the talk: 20–30 minutes, depending on how much time will be left in the course. No formal write-up is required Be ready as early as possible (presentations don't all have to be at the end of the semester, and they'll take several classes anyway).

$\longrightarrow$ Come to discuss your project with me during office hours, so I can confirm that it is appropriate and doesn't overlap with projects of other students.

Lecture 2 (Wed Sep 2)

2 Weak Lyapunov functions

2.1 LaSalle and Barbalat

Consider the general system

$$\dot{x} = f(x) \tag{7}$$

where $x \in \mathbb{R}^n$ and we assume, here and later, that $f : \mathbb{R}^n \to \mathbb{R}^n$ is sufficiently regular (at least locally Lipschitz) so that solutions of the system exist and are unique. We will typically also assume that $f(0) = 0$, to have an equilibrium at the origin.

Let V be a $\mathcal{C}^1$ (continuously differentiable) function from $\mathbb{R}^n$ to $[0, \infty)$. Such functions are usually called *candidate Lyapunov functions.* Its derivative along solutions of (7) is defined as

$$\dot{V}(x) := \frac{\partial V}{\partial x} \cdot f(x)$$

Theorem 2 *Suppose that for some nonnegative definite continuous function W we have*

$$\dot{V}(x) \leq -W(x) \leq 0 \qquad \forall\, x \tag{8}$$

Then, for every solution $x(t)$ that remains bounded for all $t \geq 0$, it is true that

$$W(x(t)) \to 0 \quad as \quad t \to \infty \tag{9}$$

Comments:

- It is very important to remember that the claim (9) is made only for bounded solutions. Boundedness of $x(t)$ will be used in the proof. So if, for example, $|x(t)| \to \infty$ then it does not necessarily satisfy (9). (There is a famous paper by Monopoli which makes this mistake while trying to prove stability of an adaptive control system, and for many years many people have been trying to determine whether his claim is still correct or not.)

- Boundedness of all solutions of (7) follows from (8) if V is radially unbounded. As we will see, sometimes boundedness of solutions can be shown by a separate argument, so we don't require it in Theorem 8.

- Note that, unlike in Theorem 1, here V is not required to be positive definite, only positive semidefinite. This is useful in adaptive control, where one often works with functions V that depend on some but not all states. Unless V is positive definite, we can't use Theorem 1 to conclude stability in the sense of Lyapunov, but we'll rarely need this.

- *LaSalle's theorem* makes a more precise claim, namely, that every bounded solution approaches the largest positive-invariant set inside the set

$$\{x : \dot{V}(x) = 0\}$$

which in turn belongs to the set $\{x : W(x) = 0\}$ in view of (8). As a corollary, we have the following stability result, known as *Barbashin-Krasovskii-LaSalle theorem*: If V is as in Theorem 2 and $\dot{V}$ does not stay identically 0 along any nonzero trajectory, then all bounded solutions converge to 0. The simpler but weaker claim of Theorem 2 remains true for time-varying systems[1] $\dot{x} = f(t, x)$, while LaSalle's theorem in general does not. For our purposes, just convergence of $W(x(t))$ to 0 will usually be enough.

PROOF OF THEOREM 2. Integrate (8) from time 0 to current time t, to get

$$V(x(t)) - V(x(0)) \leq -\int_0^t W(x(s))ds$$

or, rearranging terms,

$$\int_0^t W(x(s))ds \leq V(x(0)) - V(x(t)) \leq V(x(0)) < \infty$$

where the second inequality follows from the fact that $V(x) \geq 0 \; \forall x$. Since the above calculation is true for every t, and $V(x(0))$ is a constant (fixed by the choice of initial condition), we can take the limit as $t \to \infty$ and conclude that

$$\int_0^\infty W(x(s))ds < \infty \tag{10}$$

The integral is of course also bounded from below by 0 because W is nonnegative definite.

We now need to show that the finiteness of the above improper integral, together with the hypotheses of Theorem 2 that we haven't used yet (which ones are these?) implies (9). We formulate this remaining step as a lemma, because it is often useful as an independent result.

Lemma 3 ("Barbalat's lemma") *Let $x(t)$ be bounded and let its derivative $\dot{x}(t)$ also be bounded. Let W be a continuous function of x such that the integral*

$$\int_0^\infty W(x(s))ds$$

is well defined and finite. Then (9) *is true.*

Theorem 2 follows from Barbalat's lemma because $\dot{x} = f(x)$ and this is bounded if x is bounded.

Barbalat's lemma seems intuitively obvious. In fact, how can a function whose integral over $[0, \infty)$ is finite not converge to 0? Well, here's one possibility:

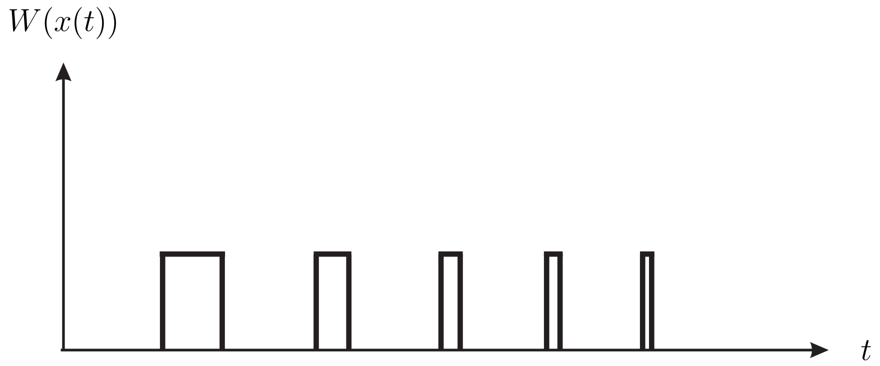

If the areas under the pulses decrease as geometric series, then the total area will be finite, even though the function does not converge to 0. Of course, we can smooth out all corners and make this function $\mathcal{C}^1$, no problem.

So, what's the catch? Why can't such behavior come from a bounded solution of the system (7)?

Reason: this requires *unbounded derivative.* Indeed, to come up to the same height and back down over shorter and shorter time intervals, the pulses would need to get steeper and steeper. So, if we invoke the condition that $\dot{x}$ is bounded, then we can indeed show rigorously that we should have convergence.

PROOF OF BARBALAT'S LEMMA. *Note:* in this proof we are still assuming for simplicity that $W \geq 0$, as in the Theorem, but with a little extra care the argument still goes through if W is not sign definite.

Suppose that $W(x(t))$ does *not* converge to 0 as $t \to \infty$. Then, there would have to exist a number $\varepsilon > 0$ and a sequence of times $\{t_k\}$ converging to ∞ at which we have

$$W(x(t_k)) \geq \varepsilon \qquad \forall\, k$$

(To see why, just recall the definition of convergence: $\forall\, \varepsilon > 0\ \exists\, T > 0$ such that $|W(x(t))| \leq \varepsilon$ $\forall\, t \geq T$.)

This situation is similar to the one in the previous figure, ε being the pulse height.

Since W is a continuous function of x, and since $x(t)$ is bounded for all t, there exists a constant $\Delta x > 0$ with the following property:

$$W(x(t)) \geq \varepsilon/2 \quad \text{whenever} \quad |x(t_k) - x(t)| \leq \Delta x$$

This is because for any choice of time t such that $x(t)$ is close enough to some $x(t_k)$, the corresponding values of W will be within $\varepsilon/2$ of each other. (This is *uniform continuity* of W as a function of x. Boundedness of $x(\cdot)$ is crucial here, because otherwise, a uniform constant Δx with the above property may not exist; think of the function $W(x) = x^2$ over an infinite range of x.)

[1]Under appropriate Lipschitzness and boundedness conditions on f; see [Khalil, Theorem 8.4].

We have a similar property for x as a function of t: there exists a constant $\Delta t > 0$ such that

$$|x(t_k) - x(t)| \leq \Delta x \quad \text{whenever} \quad |t_k - t| \leq \Delta t$$

This is true because the derivative $\dot{x}$ is bounded, so the time it takes for x to grow by Δx is bounded from above.

The function $W(x(t))$ is a composition of $W(x)$ and $x(t)$. Combining the above two properties, we see that for $t \in [t_k, t_k + \Delta t]$, $k = 1, 2, \dots$ we have $W(x(t)) \geq \varepsilon/2$. Hence, each such interval contributes an area of at least $\Delta t \cdot \varepsilon/2$ to the total integral in (10). Since there are infinitely many such intervals, the total integral cannot be finite, which is a contradiction. □

2.2 Connection with observability

(This will be useful later for understanding the persistency of excitation property.)

We would like to know under what circumstances Theorem 2 allows us to conclude asymptotic stability of the system (7), i.e., when does a weak Lyapunov function ($\dot{V} \leq 0$) let us show the same property as does a strong Lyapunov function ($\dot{V} < 0$)?

LaSalle's theorem (stated earlier) says that the largest invariant set inside the set $\{x : \dot{V}(x) = 0\}$ should be $\{0\}$. In other words, there should be no nonzero trajectories of (7) along which we have $\dot{V} \equiv 0$. For linear systems, there is a simpler way to understand this result, in terms of *observability*. This relies on Theorem 2 and so, compared to LaSalle, has the advantage that it automatically applies to time-varying systems.

Consider first a linear time-invariant (LTI) system

$$\dot{x} = Ax$$

and a quadratic candidate Lyapunov function

$$V(x) = x^T P x$$

where P is a symmetric positive definite matrix. Its derivative along solutions is given by

$$\dot{V}(x) = x^T (PA + A^T P) x$$

This is nonnegative definite if for some (not necessarily square) matrix C we have

$$PA + A^T P \leq -C^T C \leq 0 \tag{11}$$

(the second inequality is of course automatic). This gives

$$\dot{V}(x) \leq -x^T C^T C x = -y^T y$$

where we defined

$$y := Cx$$

It is convenient to think of y as a fictitious output.

All conditions of Theorem 2 are satisfied; in particular, boundedness of solutions follows because V is radially unbounded. Therefore, $y = Cx \to 0$ as $t \to \infty$.

Now, let us *assume* that (A, C) is an observable pair. This is well-known to imply that $x \to 0$ ($y \to 0 \Rightarrow x \to 0$ for observable systems, can see this from observability Gramian inversion or from output injection [ECE 515]). So, the closed-loop system must be asymptotically stable (in fact, exponentially stable since it's linear).

$\longrightarrow$ Detectability is enough to get the above implication, but we work with observability here because it is somewhat easier to define and check. We will come across detectability later, though.

The above result also extends to linear time-varying (LTV) systems, i.e., we can allow all matrices (A, P, C) to depend on time. The time-varying version of (11) is

$$\dot{P}(t) + P(t)A(t) + A^T(t)P(t) \leq -C^T(t)C(t) \leq 0 \tag{12}$$

We can get an asymptotic stability result from this, provided we have the correct observability property of the time-varying pair $(A(t), C(t))$. This goes as follows. The *observability Gramian* is defined as

$$M(t_0, t_0 + T) := \int_{t_0}^{t_0+T} \Phi^T(t, t_0)C^T(t)C(t)\Phi(t, t_0)dt$$

where $\Phi(\cdot, \cdot)$ is the system transition matrix [ECE 515]. The system is said to be *uniformly completely observable (UCO)* [Kalman, 1960] if for some positive constants T, β_1, β_2 we have

$$\beta_1 I \leq M(t_0, t_0 + T) \leq \beta_2 I \qquad \forall\, t_0$$

For UCO systems, the implication $y \to 0 \Rightarrow x \to 0$ is still true; this follows from the identity

$$\int_{t_0}^{t_0+T} \Phi^T(t, t_0)C^T(t)y(t)dt = M(t_0, t_0 + T)x(t_0)$$

For LTI systems, the dependence on t_0 disappears and UCO simply says that the observability Gramian (which is now defined more explicitly in terms of matrix exponentials) is positive definite for some choice of T. This is the usual observability notion, which is equivalent to the well-known rank condition.

We can now easily see that for LTV systems, (12) plus UCO give asymptotic stability. It can be shown that asymptotic stability is uniform (with respect to initial time), which implies that it is in fact exponential (since the system is linear). See [Khalil, Example 8.11] for an argument slightly different from the one given above, which proves the exponential stability claim rigorously. This result is also stated in [Ioannou-Sun, Theorem 3.4.8] and it will be useful for us in the sequel.

$\longrightarrow$ No lecture Monday: Labor Day

Lecture 3 (Wed Sep 9)

$\longrightarrow$ Problem Set 1 is assigned.

2.3 Back to the adaptive control example

(Modulo different notation, the same example is analyzed in Khalil, p. 130.)

In Example 1 we had the closed-loop system

$$\dot{x} = (\theta - \hat{\theta} - 1)x \tag{13}$$

$$\dot{\hat{\theta}} = x^2 \tag{14}$$

and the candidate Lyapunov function

$$V(x, \hat{\theta}) = \frac{x^2}{2} + \frac{(\hat{\theta} - \theta)^2}{2}$$

with derivative

$$\dot{V}(x, \hat{\theta}) = (\theta - \hat{\theta} - 1)x^2 + (\hat{\theta} - \theta)x^2 = -x^2 \leq 0$$

Let's check carefully that Theorem 2 indeed applies. V is radially unbounded, hence all trajectories are bounded (as we said earlier). Define

$$W(x, \hat{\theta}) := x^2$$

which is nonnegative definite and continuous. (Remember that the system state includes both x and $\hat{\theta}$.) Thus by Theorem 2 we have $W(x(t)) \to 0$, hence $x(t)$ converges to 0 as needed.

On the other hand, Theorem 2 tells us nothing about $\hat{\theta}(t)$. It may not converge to θ, or to anything else. Indeed, the line in the $(x, \hat{\theta})$ space given by

$$\{(x, \hat{\theta}) : x = 0\}$$

consists entirely of equilibria (this is obvious from the system equations). This line is exactly the set

$$\{(x, \hat{\theta}) : \dot{V}(x, \hat{\theta}) = 0\}$$

This set is invariant, so even the stronger LaSalle's theorem gives nothing further than convergence to this line. For example, it is clear that if we start on this line, with any value of $\hat{\theta}$, then we'll stay there and $\hat{\theta}$ won't change.

$\longrightarrow$ Since the value of θ is unknown, an interesting feature of the above Lyapunov function is that it's not completely known, but is instead an abstract function whose existence is guaranteed. This is typical in adaptive control.

When we first studied the example, we also tried the candidate Lyapunov function

$$V(x, \hat{\theta}) = \frac{x^2}{2} \tag{15}$$

whose derivative is

$$\dot{V} = (\theta - \hat{\theta} - 1)x^2 \tag{16}$$

This can be both positive or negative, so Theorem 2 does not apply. In fact, we can't even use this V to show boundedness of solutions. However, with a bit of clever analysis and by using Barbalat's lemma, we can still use this V to show that $x(t) \to 0$ as follows.

First, use the $\dot{\hat{\theta}}$ equation (14) to rewrite (16) as

$$\dot{V} = (\theta - \hat{\theta} - 1)\dot{\hat{\theta}} = (\theta - 1)\dot{\hat{\theta}} - \hat{\theta}\dot{\hat{\theta}}$$

Integrate this to get

$$\frac{x^2(t)}{2} = (\theta - 1)\hat{\theta}(t) - \frac{1}{2}\hat{\theta}^2(t) + C \tag{17}$$

where C is a constant determined by the initial conditions (of x and $\hat{\theta}$).

We know from (14) that $\hat{\theta}$ is monotonically nondecreasing. Thus it must either approach a finite limit or grow without bound. But if it were unbounded, then the right-hand side of (17) would become negative, while the left-hand side cannot be negative—a contradiction. This shows that $\hat{\theta}(t)$ is bounded. Looking at (17) again, we see that $x(t)$ must then also be bounded.

Integrating (14), we have

$$\int_0^t x^2(s)ds = \hat{\theta}(t)$$

We just showed that this is bounded for all t, so in the limit as $t \to \infty$ we have

$$\int_0^\infty x^2(s)ds < \infty$$

We already know that x is bounded, and $\dot{x}$ is also bounded because of (13). So, we can apply Barbalat's lemma and conclude that $x(t) \to 0$ as $t \to \infty$, and we're done.

This argument is certainly more complicated than the previous one based on Theorem 2. However, it is often difficult to find a weak Lyapunov function satisfying the hypotheses of Theorem 2, and so it is useful to have as many tools as possible at our disposal. The above steps—integrating differential equations, showing boundedness, and invoking Barbalat's lemma—are used very often in adaptive control.

3 Minimum-phase systems and universal regulators

The control law in the above example is capable of stabilizing any plant from the class of scalar plants

$$\dot{x} = \theta x + u$$

parameterized by $\theta \in \mathbb{R}$.

A controller is called a *universal regulator* for a given class of plants if, when connected to *any* plant from this class, it guarantees that:

- All signals in the closed-loop system remain bounded, and
- The plant's state converges to 0 as $t \to \infty$

under arbitrary initial conditions (of both the plant and the controller).

Note: even though we used the notation $\hat{\theta}$, we are not really trying to estimate the value of θ. We have no reason to believe that $\hat{\theta} - \theta$ becomes small in any sense whatsoever. We can think of a universal regulator as doing some kind of *exhaustive search* through the space of possible controller gains, so that eventually it will find the gains that will stabilize the given unknown plant, whatever it is (as long as it is in a prescribed class). Later, we will study different controllers which explicitly rely on parameter estimation and certainty equivalence. (See Section 6.5 and Section 8.2.)

$\longrightarrow$ It is hard to make a rigorous distinction between universal regulators and estimation-based adaptive controllers. This is because in both cases the controller is dynamics, and it's impossible to define formally when these dynamics are estimating anything and when they are not. The difference is not so much in the appearance of the controller, but in the *design philosophy* behind it. With some experience, it is usually easy to tell one from another.

Exhaustive search means that transient behavior might be poor (and this is true). However, from the theoretical point of view the possibility of designing universal regulators for large classes of plants is quite interesting.

Next, we want to see if universal regulators exist for larger classes of plants.

3.1 Universal regulators for scalar plants

We study the class of scalar plants

$$\dot{y} = ay + bu \tag{18}$$

parameterized by a and b, where a is completely unknown and $b \neq 0$.

We write y, not x, because later we'll generalize to higher-dimensional plants with scalar outputs.

3.1.1 The case $b > 0$

We first consider the simpler case where the sign of the control gain b (the "high-frequency gain") is known, and with no loss of generality we take it to be positive.

It turns out that this makes relatively minor difference with Example 1 (where b was equal to 1). In fact, the same controller works, and almost the same analysis applies.

Controller:

$$\begin{aligned} \dot{k} &= y^2 \\ u &= -ky \end{aligned} \tag{19}$$

This notation better reflects the fact that we're not estimating a and b, but just searching the space of controller gains for a suitable one. The differential equation for k is a *tuning law*, or *tuner*. (k here plays the role of $\hat{\theta} + 1$ before.)

Closed-loop plant:

$$\dot{y} = (a - bk)y \tag{20}$$

Candidate Lyapunov function—?

Recall that in Example 1 we examined two different choices. One was (3), with which the analysis was simple via Theorem 2. The other choice was (15), which itself is simpler but the analysis was a bit more complicated. For the present case of $b > 0$, both choices would still work. However, for the general case of b with arbitrary sign, only the second one will be helpful. So, we go with

$$V := \frac{y^2}{2}$$

Its derivative is

$$\dot{V} = (a - bk)y^2 = (a - bk)\dot{k}$$

Integrate this to get

$$\frac{y^2(t)}{2} = ak(t) - \frac{bk^2(t)}{2} + C \tag{21}$$

where C is a constant determined by the initial conditions.

We know from the tuning law (19) that k is monotonically nondecreasing. Thus it must either approach a finite limit or grow without bound. But if it were unbounded, then the right-hand side of (21) would become negative (since $b > 0$; that's where we're using this assumption), while the left-hand side cannot be negative—a contradiction. Hence k is bounded, and so y is also bounded.

Integrating (19), we have

$$\int_0^t y^2(s)ds = k(t)$$

Since k is bounded, this means $y \in L^2$.

We already know that y is bounded, and $\dot{y}$ is also bounded because of (20). By Barbalat, $y(t) \to 0$ as $t \to \infty$.

Our conclusion so far is that the above controller is a universal regulator for the class of plants given by (18) for all $a \in \mathbb{R}$ and $b > 0$.

3.1.2 General case: non-existence results

We now go back to the general case where the sign of b in (18) can be arbitrary (as long as $b \neq 0$). It turns out that this extra uncertainty makes the problem of designing a universal regulator significantly more challenging.

Let us consider a controller of the form

$$\begin{aligned} \dot{z} &= f(z, y) \\ u &= h(z, y) \end{aligned} \tag{22}$$

where $z \in \mathbb{R}$ (i.e., the controller dynamics are scalar) and f and h are continuous *rational* functions (i.e., ratios of polynomials with no real poles).

$\longrightarrow$ Our previous controller is of this form—it's polynomial.

Claim: No such 1-D rational controller can be a universal stabilizer for (18) with unknown sign of b.

Closed-loop system:

$$\begin{aligned} \dot{y} &= ay + bh(z, y) \\ \dot{z} &= f(z, y) \end{aligned} \tag{23}$$

$\longrightarrow$ Remember the definition of a universal regulator: it must globally stabilize all plants in the class. So, to show that the above controller *cannot* be a universal regulator, we must show that for every choice of f and h, there exists a choice of values for a and b such that $y \not\to 0$ (at least for some special bad choices of initial conditions).

First, we show that f cannot be an identically zero function. If it were, then the value of z would be constant: $z \equiv z_0$, and the plant would be

$$\dot{y} = ay + bh(z_0, y)$$

(the controller would then be static).

Suppose first that $h(z_0, y) \neq 0 \ \forall\, y > 0$. Then it maintains the same sign. Assume that $h(z_0, y) > 0 \ \forall\, y > 0$ (the other case is completely analogous). Choose, for example,

$$a = 1, \qquad b = 1$$

Then we have

$$\dot{y} = y + h(z_0, y) > 0 \qquad \forall\, y > 0$$

This means that for a positive initial condition, the value of y can only increase, so it cannot go to 0.

Now suppose that $h(z_0, y_0) = 0$ for some $y_0 > 0$. Then take, for example,

$$a = 0, \qquad b \text{ arbitrary}$$

This gives

$$\dot{y} = bh(z_0, y)$$

for which $y = y_0$ is an equilibrium. Again, y cannot go to 0 from all initial conditions.

So, f is not identically zero, i.e., the controller has nontrivial dynamics.

$\longrightarrow$ Note: in the above argument, rationality of h was in fact not used. In other words, *no* static controller can be a universal regulator. (Rationality will become important when f is not identically zero, see below.)

Lecture 4 (Mon Sep 14)

Therefore, there must exist a z_0 for which $f(z_0, y)$ is a nonzero function of y. Since it is rational, it has finitely many zeros. Hence, there exists a $y_0 > 0$ such that $f(z_0, y)$ has the same sign for all $y \geq y_0$. Assume that this sign is positive (the other case is similar):

$$f(z_0, y) > 0 \qquad \forall\, y \geq y_0 \tag{24}$$

Now consider $h(z, y_0)$ as a function of z. It is also rational, hence there exists a $z_1 \geq z_0$ such that $h(z, y_0)$ has the same sign for all $z \geq z_1$. Again, assume for concreteness that this sign is positive.

By continuity, $h(z, y_0)$ is then *bounded from below* for $z \geq z_0$ (by some possibly negative but finite number).

Now, pick $a > 0$, and then pick $b > 0$ small enough so that

$$ay_0 + bh(z, y_0) > 0 \qquad \forall z \geq z_0 \tag{25}$$

Let's now look at the plant (23) and the inequalities (24) and (25). We see that solutions cannot leave the region

$$\{(y, z) : y \geq y_0,\ z \geq z_0\}$$

because everywhere on the boundary of this region they are directed inside. In other words, it is an *invariant region* for the closed-loop system.

Therefore, convergence of y to 0 (i.e., convergence to the z-axis) from initial conditions inside this region is not achieved.

Working out the remaining cases is left as an exercise.

3.1.3 Nussbaum gains

Does a universal regulator exist at all in the general case of unknown sign of b?

R. Nussbaum showed in 1983 that it does. (It is interesting that he's a pure mathematician at Rutgers and doesn't work in adaptive control; he learned about this problem from a paper by Morse and came up with the idea. His paper is in Systems and Control Letters, vol. 3, pp. 243-246; the solution we give here is taken from Morse's notes and is a bit simpler than what Nussbaum had. The previous non-existence result is also established by Nussbaum in the same paper.)

We leave the tuner as before:

$$\dot{k} = y^2$$

but let the control law be

$$u = -N(k)ky$$

(whether or not to put a minus sign in front of N above is an arbitrary choice), where $N(\cdot)$ is some function satisfying

$$\sup_{k>0} \frac{1}{k} \int_0^k N(s)s ds = \infty,$$
$$\inf_{k>0} \frac{1}{k} \int_0^k N(s)s ds = -\infty$$

In other words, this function keeps crossing back and forth between positive and negative values while giving higher and higher absolute value of the above integral.

Example: $N(k) = k \cos k$.

$\longrightarrow$ Since N must clearly have an infinite number of zeros, it is not rational, so the previous non-existence result does not apply to this controller.

With this modified controller, the analysis goes through much like before:

$$V = \frac{y^2}{2}$$

gives

$$\dot{V} = (a - bN(k)k)y^2 = (a - bN(k)k)\dot{k}$$

Integrating, we get

$$\frac{y^2(t)}{2} = ak(t) - b \int_0^{k(t)} N(s)s ds + C \tag{26}$$

where C is a constant determined by the initial conditions.

We have as before that k is monotonically nondecreasing, so it either has a finite limit or grows to infinity. If it were unbounded, then the right-hand side of (26) would eventually become negative (no matter what the sign of b is!) Indeed, this will happen when k reaches a value $\bar{k}$ for which

$$b\frac{1}{\bar{k}} \int_0^{\bar{k}} N(s)s ds > a + \frac{C}{\bar{k}}$$

which exists by the defining properties of $N(\cdot)$.

But the left-hand side cannot be negative, and we reach a contradiction. Hence k is bounded, and so y is also bounded.

After this, the analysis is exactly the same as before: $y \in L^2$ (since k is bounded), y is bounded, $\dot{y}$ is bounded, and by Barbalat $y \to 0$.

An important observation is that in the above construction, $N(\cdot)$ doesn't need to be continuous. We do need it to be locally bounded (to ensure that boundedness of k implies boundedness of $N(k)$

and hence boundedness of $\dot{y}$) and nice enough so that solutions of the closed-loop system are well defined. Piecewise continuous $N(\cdot)$ will do.

In fact, *piecewise constant* $N(\cdot)$ is probably the easiest to design.

The values of N don't actually need to grow in magnitude, they can just be maintained on longer and longer intervals to satisfy the condition. This is an example of a *switching logic*.

This design was proposed by Willems and Byrnes in 1984. It was probably the first example of a switching adaptive control algorithm, which motivated a lot of subsequent work. We'll talk more about switching adaptive control later in the course (although in a different context: estimator-based).

3.2 Relative degree and minimum phase

To move beyond scalar plants, we need a few new concepts.

Consider a linear system

$$\begin{aligned}\dot{x} &= Ax + Bu \\ y &= Cx\end{aligned}$$

Assume that it is *SISO*, i.e., u and y are scalar (and so B and C are vectors, not matrices). We'll briefly mention MIMO case later.

Relative degree is, basically, the number of times that we need to differentiate the output until the input appears.

$$\dot{y} = CAx + CBu$$

where CB is actually a scalar since the system is SISO.

If $CB \neq 0$, then relative degree equals 1. If $CB = 0$, then we differentiate again:

$$\ddot{y} = CA^2x + CABu$$

and so on. Relative degree is the positive integer r such that

$$CB = CAB = \cdots = CA^{r-2}B = 0, \qquad CA^{r-1}Bu \neq 0$$

which means that $y^{(r)}$ is the first derivative that depends on u because

$$y^{(r)} = CA^r x + CA^{r-1}Bu$$

The terminology "relative degree" is motivated by the following. Consider the transfer function of the system:

$$G(s) = C(Is - A)^{-1}B = \frac{q(s)}{p(s)} \tag{27}$$

Then

$$r = \deg(p) - \deg(q)$$

i.e., relative degree is the difference between the degrees of the denominator and numerator polynomials.

To define relative degree for MIMO systems, we need to differentiate *each* output until *some* input appears. We'll have (assuming both y and u have the same dimension m)

$$\begin{pmatrix} y_1^{(r_1)} \\ y_2^{(r_2)} \\ \vdots \\ y_m^{(r_m)} \end{pmatrix} = Lx + Mu$$

The matrix M generalizes $CA^{r-1}B$ we had earlier. If M is nonsingular, then $(r_1, \dots, r_m)$ is called the *vector relative degree*. It may not exist (the matrix M may be singular).

The concept of relative degree also extends quite easily to nonlinear systems. We'll only deal with nonlinear systems that are SISO and affine in controls.

Lecture 5 (Wed Sep 16)

Example 2

$$\begin{aligned}\dot{x}_1 &= x_3 - x_2^3\\ \dot{x}_2 &= -x_2 - u\\ \dot{x}_3 &= x_1^2 - x_3 + u\\ y &= x_1\end{aligned}$$

Differentiate y:

$$\dot{y} = x_3 - x_2^3$$

and this doesn't depend on u. Differentiate again:

$$\ddot{y} = x_1^2 - x_3 + u - 3x_2^2(-x_2 - u) = x_1^2 - x_3 + 3x_2^3 + (1 + 3x_2^2)u$$

Since

$$1 + 3x_2^2 \neq 0$$

relative degree is $r = 2$.

$\longrightarrow$ In this example, $r = 2$ globally in the state space. In general for nonlinear systems, relative degree is a local concept, because the term multiplying u may vanish for some x. (For linear systems this doesn't happen because it is a constant.)

Defining new coordinates

$$z_1 := x_1 = y, \qquad z_2 := x_3 - x_2^3$$

we get

$$\begin{aligned}\dot{z}_1 &= z_2\\ \dot{z}_2 &= x_1^2 - x_3 + 3x_2^3 + (1 + 3x_2^2)u\end{aligned}$$

We want to complete a coordinate transformation and write everything in terms of z. For this we need z_3. We can always find z_3 whose differential equation doesn't depend on u. In this example, we can take

$$z_3 := x_2 + x_3$$

for which

$$\dot{z}_3 = x_1^2 - x_2 - x_3 = z_1^2 - z_3$$

and this doesn't depend on u.

Note that the Jacobian of the map from x to z is

$$\begin{pmatrix}1 & 0 & 0\\ 0 & -3x_2^2 & 1\\ 0 & 1 & 1\end{pmatrix}$$

which is nonsingular, so the coordinate transformation is well defined (in fact, globally).

Can also check that this transformation preserves the origin.

The z-dynamics are of the form

$$\begin{aligned}
\dot{z}_1 &= z_2 \\
\dot{z}_2 &= b(z_1, z_2, z_3) + a(z_1, z_2, z_3)u \\
\dot{z}_3 &= z_1^2 - z_3 \\
y &= z_1
\end{aligned} \tag{28}$$

where we know that

$$a(z_1, z_2, z_3) \neq 0 \qquad \forall\, z$$

This system is in so-called *normal form.* It has a chain of integrators of length $r-1$ (where r is relative degree), followed by a state-dependent affine function of u with the term multiplying u being non-zero, and completed by the rest of the dynamics that do not involve u. □

The next concept we need is that of a *minimum-phase* system. A linear SISO system is called minimum-phase if its *zeros* have negative real parts. These are the roots of the numerator polynomial $q(s)$ in the transfer function (27).

(The term "minimum-phase" comes from the fact that among transfer functions with the same magnitude Bode plot, the one with stable zeros has the smallest variation of the phase Bode plot.)

The interpretation of this notion is that the "inverse system" is asymptotically stable. For example, the inverse of

$$y = \frac{s+2}{(s+3)^2}u$$

is

$$u = \frac{(s+3)^2}{s+2}y$$

The problem is that this is not proper. But we can fix this if we work with derivatives of y instead. For example, we can multiply both sides of the original equation by $s+1$:

$$\dot{y} + y = \frac{(s+1)(s+2)}{(s+3)^2}u$$

This has relative degree 0 and so we can invert it without losing properness:

$$u = \frac{(s+3)^2}{(s+1)(s+2)}(\dot{y} + y)$$

Since the original system has a stable zero (at $s = -2$), and since we were careful not to add an unstable zero, the inverse system has stable poles (which are the same as the zeros of the initial system). Asymptotic stability of this inverse system implies, in particular:

$$y \equiv 0 \;\Rightarrow\; x, u \to 0$$

where x is the internal state in any minimal (controllable and observable) state-space realization.

This last implication is a good way of thinking about minimum-phase. Suppose we chose a control u which maintains the output to be identically zero ($y \equiv 0$). The resulting dynamics—called *zero dynamics*—should then be asymptotically stable ($x \to 0$, and consequently $u \to 0$ because it is an output of the inverse system).

The minimum-phase property (as well as the relative degree) are coordinate-independent, but they are easier to study using a normal form. Consider again the normal form (28). We want $y = z_1 \equiv 0$. Then we should also have $\dot{y} = z_2 \equiv 0$. This in turn requires $\dot{z}_2 \equiv 0$. Since the differential equation for z_2 is controllable ($a \neq 0$), we can apply the feedback control law

$$u = -\frac{b(z_1, z_2, z_3)}{a(z_1, z_2, z_3)}$$

to enforce this.

The zero dynamics are the remaining dynamics of the system, constrained by the condition $y \equiv 0$. I.e., zero dynamics describe the remaining freedom of motion after we fix the initial conditions (in the present case, $z_1(0) = z_2(0) = 0$) and select the control as above.

Here, the zero dynamics are

$$\dot{z}_3 = -z_3$$

So they are 1-dimensional and asymptotically stable. Hence, this system is minimum-phase. (For nonlinear systems one *defines* minimum-phase via asymptotic stability of zero dynamics, since there is no transfer function.)

Can check that $u \to 0$ because b is 0 at 0.

3.3 Universal regulators for higher-dimensional plants

So far we've discussed universal regulators of the scalar plant (18), which was

$$\dot{y} = ay + bu$$

We now want to consider higher-dimensional linear plants. From now on y will be an output; it'll still be scalar.

To get a class of systems which is large enough but for which the universal regulator problem is still manageable, we assume that:

- The plant is SISO.
- The plant has relative degree 1.
- The plant is minimum-phase.

From the earlier discussion about normal forms it is clear that, up to a coordinate transformation, we can represent all such plants in the form

$$\begin{aligned}\dot{y} &= ay + bu + c^T z\\ \dot{z} &= Az + dy\end{aligned}$$

where A is a *Hurwitz* matrix (because it gives zero dynamics), c and d are vectors, a and $b \neq 0$ are scalars, and z is a vector which together with y gives the plant state: $x = (y, z)$. The entries of A, a, b, c, d are all *unknown.*

The goal is to have $y, z \to 0$ as $t \to \infty$ while keeping all closed-signals bounded. And we want to achieve this with output feedback (using only y, not z).

This class of plants includes our previous scalar case, but it's obviously much larger. Nevertheless, it turns out that the *same* controller as in Section 3.1.3: the tuner

$$\dot{k} = y^2$$

and the controller

$$u = -N(k)ky$$

(where N is a Nussbaum gain function) solves the regulation problem for this class of plants!

The intuition here is that in some sense, minimum-phase systems with relative degree 1 behave essentially like scalar systems: if you stabilize y, the rest of the state will also be automatically stabilized. The analysis needed to show this becomes more challenging, though.

The closed-loop system is

$$\begin{aligned}\dot{y} &= (a - bN(k)k)y + c^T z\\ \dot{z} &= Az + dy\\ \dot{k} &= y^2\end{aligned}$$

As in the scalar case, let's try the Lyapunov function

$$V = \frac{y^2}{2}$$

It gives

$$\dot{V} = (a - bN(k)k)y^2 + yc^T z = (a - bN(k)k)\dot{k} + yc^T z$$

Integrating, we get

$$\frac{y^2(t)}{2} = ak(t) - b\int_0^{k(t)} N(s)sds + \underbrace{\int_0^t y(s)c^T z(s)ds}_{\text{new term}} + C \tag{29}$$

where C is a constant determined by the initial conditions.

Lecture 6 (Mon Sep 21)

We can view $c^T z$ as an output of the exponentially stable linear system

$$\dot{z} = Az + dy \tag{30}$$

In what follows, we'll need some well-known facts about linear systems. Proofs of these facts will either be assigned as homework or are variations on those to be assigned in homework.

FACT 1: Exponentially stable linear systems have finite L_2 induced norms.

This means, in our case, that there exist constants c_1, c_2 such that

$$\int_0^t (c^T z(s))^2 ds \le c_1 |z(0)|^2 + c_2 \int_0^t y^2(s) ds$$

From this, the Cauchy-Schwartz inequality, and simple square completion we have

$$\begin{aligned}
\int_0^t y(s) c^T z(s) ds &\le \sqrt{\int_0^t y^2(s) ds} \sqrt{\int_0^t (c^T z(s))^2 ds} \\
&\le \frac{1}{2} \int_0^t y^2(s) ds + \frac{1}{2} \int_0^t (c^T z(s))^2 ds \\
&\le c_3 |z(0)|^2 + c_4 \int_0^t y^2(s) ds = c_3 |z(0)|^2 + c_4 (k(t) - k(0))
\end{aligned}$$

where

$$c_3 = \frac{c_1}{2}, \qquad c_4 = \frac{c_2}{2} + \frac{1}{2}$$

Substituting this bound into (29), we get

$$\frac{y^2(t)}{2} \le (a + c_4) k(t) - b \int_0^{k(t)} N(s) s ds + \bar{C}$$

Now that we were able to handle the extra term, we're almost done. Exactly as before, we show that k and y are bounded, and $y \in L_2$.

How do we show that $\dot{y}$ is bounded? Need to know that z is bounded.

FACT 2: Given an exponentially stable linear system, if the input is bounded, then the state is bounded.

Applying this to the system (30) and using boundedness of y, we have that z is bounded.

$y \in L_2$, $\dot{y}$ is bounded $\implies$ $y \to 0$ (Barbalat).

FACT 3: Given an exponentially stable linear system, if the input is in L_2 or converges to 0, then the state converges to 0.

From this we conclude that $z \to 0$. Thus the regulation objective is fulfilled.

Designing universal regulators for systems with relative degree 2 or higher is more challenging. In Section 5 we will study *backstepping*, which a tool for overcoming the relative degree obstacle—but we won't be working in the context of universal regulators any more.

Additional reading on universal regulators: A. Ilchmann, "Non-identifier based adaptive control of dynamical systems: A survey," IMA J. Math. Contr. Inform., vol. 8, pp. 321–366, 1991.

4 Lyapunov-based design

In Example 1, we started with the system

$$\dot{x} = \theta x + u$$

How did we come up with the adaptive control design?

The feedback law

$$u = -(\hat{\theta} + 1)x \tag{31}$$

was motivated by the certainty equivalence principle. This makes sense, but since parameter estimates don't converge, it is *not* a rigorous justification why this controller is stabilizing.

Then we had the tuning law

$$\dot{\hat{\theta}} = x^2 \tag{32}$$

for which we didn't really have any justification, except to say that the quadratic dependence of the right-hand side on x should dominate linear growth of the plant.

Then we took the candidate Lyapunov function

$$V(x, \hat{\theta}) = \frac{x^2}{2} + \frac{(\hat{\theta} - \theta)^2}{2} \tag{33}$$

and showed that its derivative satisfies

$$\dot{V}(x, \hat{\theta}) = -x^2 \tag{34}$$

from which, by Theorem 2, we have $x(t) \to 0$. The choice of V was, again, not systematic and involved trial and error.

If we're going to base stability analysis on a (weak) Lyapunov function, then an alternative approach is to *start* by picking V as in (33), and then design the control law and tuning law so as to get (34). The step of choosing V still involves guessing, but then V provides the basis for controller design and we get stability by construction.

For the above V, but keeping u and $\dot{\hat{\theta}}$ unspecified for now, we get

$$\dot{V} = x(\theta x + u) + (\hat{\theta} - \theta)\dot{\hat{\theta}} = xu + \hat{\theta}\dot{\hat{\theta}} + \theta(x^2 - \dot{\hat{\theta}})$$

The last term is not going to give us anything useful, because θ is the unknown parameter and we don't know its sign. So, it makes sense to cancel it. This immediately suggests the tuning law (32). We now have

$$\dot{V} = xu + \hat{\theta}x^2$$

and so, if we want (34), we need to pick u such that

$$xu + \hat{\theta}x^2 = -x^2$$

It is clear that the control law (31) is, in fact, a unique control law that gives this.

So, Lyapunov-based design allows us to reconstruct our original choices in a more methodical way. Also note that it gives us more flexibility. For example, we see that any u for which

$$xu + \hat{\theta}x^2 \leq -x^2$$

would give us the same conclusion (since strict equality is not required in Theorem 2). So, for example, any control of the form

$$u = -(\hat{\theta} + k)x, \qquad k > 1$$

works just as well. This may be somewhat obvious for this example, but in more complicated situations Lyapunov-based design can make it easier to see which controllers are stabilizing.

4.1 Control Lyapunov functions

A general concept that formalizes the idea of Lyapunov-based control design is that of a *control Lyapunov function*, or CLF. It was introduced by Artstein in 1983. Among the texts for this course, it is discussed in KKK book, Section 2.1.2.

Consider a general control system

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

where $x \in \mathbb{R}^n$ is the state and $u \in \mathbb{R}^m$ is the control input. Let V be a function having the usual properties of a candidate Lyapunov function, i.e., it is $\mathcal{C}^1$ and at least nonnegative definite.

$\longrightarrow$ In fact, here for simplicity we will work with the more common situation of a strong Lyapunov function described by the condition (6) in Theorem 1, and not a weak Lyapunov function as in Theorem 2. So, we assume that V is strictly positive definite, and that our goal is to make $\dot{V}$ strictly negative (by choice of u).

The scenario with nonstrict inequalities (weak CLF) can be developed similarly. We will see an example of this.

Let's say that V is a CLF for our system if for each $x \neq 0$ there is a value of u for which

$$\dot{V}(x,u) := \frac{\partial V}{\partial x} \cdot f(x,u) < 0$$

Somewhat more formally, but equivalently:

$$\inf_u \left\{ \frac{\partial V}{\partial x} \cdot f(x,u) \right\} < 0 \qquad \forall\, x \neq 0$$

Now, suppose that we're interested in finding a *continuous feedback law*

$$u = k(x)$$

which makes the closed-loop system globally asymptotically stable. And we want our V to be a Lyapunov function for the closed loop, i.e., we want

$$\frac{\partial V}{\partial x} \cdot f(x,k(x)) < 0 \qquad \forall\, x \neq 0$$

In fact, for the closed-loop system to be well-posed, we probably need more regularity from the feedback law—at least local Lipschitzness—but continuity is the absolute minimum we should ask for.[2] It is often also required that $k(0) = 0$, to preserve the equilibrium at the origin (assuming also that $f(0,0) = 0$).

Does the existence of a CLF imply the existence of such a stabilizing feedback law?

One is tempted to say yes. However, being able to find a value of u for each x does not automatically imply being able to glue them together into a continuous function $k(x)$. This is known as the *continuous selection problem*: from a given collection of sets parameterized by x (in our case, the sets of "good" values of u) select a continuous function of x.

[2]Continuity of the right-hand side of an ODE is enough for existence of solutions, but not enough for their uniqueness (see [Khalil, p. 88] for a reference).

Lecture 7 (Wed Sep 23)

COUNTEREXAMPLE: [Sontag and Sussmann, 1980]

$$\dot{x} = x\left[(u-1)^2 - (x-1)\right]\left[(u+1)^2 + (x-2)\right]$$

Let

$$V(x) := \frac{x^2}{2}$$

then

$$\dot{V} = x^2\left[(u-1)^2 - (x-1)\right]\left[(u+1)^2 + (x-2)\right]$$

For this to be negative, one (and only one) of the expressions in square brackets must be negative. It is easy to see that the points in the (x, u) plane where this happens are given by the interiors of the two parabolas in the picture.

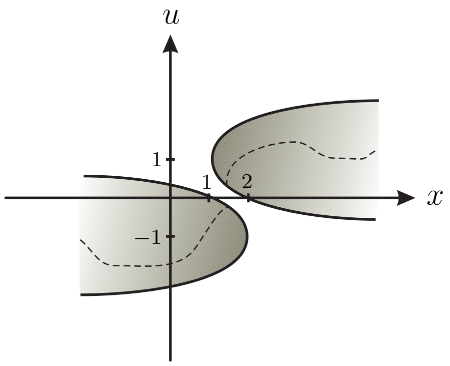

The projection of the union of the two parabolas onto the x-axis covers the whole axis. This means that V is a CLF (directly from the definition). On the other hand, the parabolas do not intersect, which means that no continuous feedback law exists that makes $\dot{V}$ negative. Any such feedback law would have to pass somehow from one parabola to the other, see the dashed curve in the picture.

$\longrightarrow$ There are several possibilities for overcoming this difficulty. One is to consider discontinuous feedback laws (or time-varying ones, or switching ones). Another is to look for classes of systems for which continuous feedback laws can be constructed.

For now, let's stick with continuous feedback laws. For *systems affine in controls*, all is well. Not only do continuous stabilizers exist [Artstein, 1983], but they can be generated from a CLF by an explicit formula [Sontag, 1989]. However, it is clear from the preceding discussion that it will take some effort to establish this result (it is far from trivial, and the affine structure must somehow play a role).

Consider the system

$$\dot{x} = f(x) + G(x)u = f(x) + \sum_{i=1}^{m} g_i(x)u_i$$

Here f and g_i are n-vectors, and G is a matrix whose columns are the g_i's. We assume that $f(0) = 0$.

The definition of CLF in this case becomes

$$\inf_u \left\{ \frac{\partial V}{\partial x} \cdot f(x) + \sum_{i=1}^{m} \frac{\partial V}{\partial x} \cdot g_i(x)u_i \right\} < 0 \qquad \forall\, x \neq 0$$

It is easy to see that this is equivalent to the condition that for all $x \neq 0$,

$$\frac{\partial V}{\partial x} \cdot g_i(x) = 0 \;\; \forall\, i \qquad \Longrightarrow \qquad \frac{\partial V}{\partial x} \cdot f(x) < 0$$

Indeed, since the controls are unbounded, we can always pick u to get

$$\frac{\partial V}{\partial x} \cdot f(x) + \sum_{i=1}^{m} \frac{\partial V}{\partial x} \cdot g_i(x)u_i < 0 \tag{35}$$

except for those x at which the terms with u are all 0, where we lose control authority and need the first term to be negative by itself.

To simplify notation, define

$$a(x) := \frac{\partial V}{\partial x} \cdot f(x), \qquad b(x) := \left(\frac{\partial V}{\partial x} \cdot g_1(x), \dots, \frac{\partial V}{\partial x} \cdot g_m(x) \right)^T$$

In particular, the CLF property can be written as

$$|b(x)| = 0 \qquad \Longrightarrow \qquad a(x) < 0 \tag{36}$$

for all $x \neq 0$.

Consider the feedback law

$$k(x) = K(a(x), b(x)) := \begin{cases} -\dfrac{a + \sqrt{a^2 + |b|^4}}{|b|^2} b, & b \neq 0 \\ 0, & b = 0 \end{cases} \tag{37}$$

It can be shown that the apparent singularity at $b = 0$ is removable. In fact, K is *analytic* as a function of a and b, except at $x = 0$ (where (36), which is used for showing this, does not hold). This means that the above control law does not lead to any loss of smoothness in closed loop; for example, it is smooth (away from $x \neq 0$) if f, G, and V are smooth.

It is also not hard to show that this feedback stabilizes the closed-loop system, with V as Lyapunov function. We do this now. Write

$$\begin{aligned}\frac{\partial V}{\partial x}\cdot\Big(f(x)+\sum_{i=1}^m g_i(x)u_i\Big) &= a(x)+\sum_{i=1}^m b_i(x)u_i\\ &= a-\sum_{i=1}^m b_i\frac{a+\sqrt{a^2+|b|^4}}{|b|^2}b_i = a-\frac{a+\sqrt{a^2+|b|^4}}{|b|^2}\sum_{i=1}^m b_i^2\\ &= -\sqrt{a^2+|b|^4}<0 \qquad \forall\, x\neq 0\end{aligned}$$

where the very last inequality follows from (36). The claim follows from Theorem 1.

The reason to put $|b|^4$ and not, e.g., $|b|^2$ inside the square root is to ensure the above smoothness property of the control law, as well as its continuity at 0 under an additional hypothesis (as discussed below).

Note that formally, to talk about global asymptotic stability of the zero equilibrium we need to make sure that $x = 0$ is indeed an equilibrium. This is why we need $f(0) = 0$ (since $k(0) = 0$ by construction).

As for $x = 0$, the feedback law (37) is automatically continuous there if V has the property that for small x, the values of u that give (35) can also be chosen small (*small control property*). This is not always possible. For example, the scalar system

$$\dot{x} = x + x^2 u$$

cannot be stabilized with a feedback that is continuous at 0, because x^2 is small near 0 compared to x so u needs to be large there, and u needs to be negative for $x > 0$ and positive for $x < 0$. The function $V(x) = x^2/2$ is a CLF for this system, but it doesn't have the small control property. If x and x^2 were flipped, then all would be well.

Anyway, continuity of $k(x)$ at 0 is not so crucial: if it is continuous away from 0, then the closed-loop system is well-posed for $x(0) \neq 0$ and all solutions go to 0, which is what we want.

The formula (37) is known as *Sontag's formula*, or *universal formula*. Similar formulas exist for control spaces different from $\mathbb{R}^m$ (e.g., bounded controls).

Example 3

$$\dot{x} = -x^3 + u, \qquad x, u \in \mathbb{R}$$

Even without using Sontag's formula, there are several rather obvious stabilizing feedback laws we can apply. For example:

$$u = x^3 - x \quad \Longrightarrow \quad \dot{x} = -x$$

This an example of *feedback linearization* design: cancel nonlinearities and get a stable linear closed-loop system.

However, the previous feedback law requires very large control effort for large x, while $-x^3$ is actually a "friendly" nonlinearity which we don't need to cancel. Indeed, consider

$$u = -x \qquad \Longrightarrow \qquad \dot{x} = -x^3 - x$$

This is globally asymptotically stable:

$$V(x) = \frac{x^2}{2}$$

gives

$$\dot{V} = -x^4 - x^2 < 0 \qquad \forall\, x \neq 0$$

A third option is to do nothing:

$$u = 0 \qquad \Longrightarrow \qquad \dot{x} = -x^3$$

and the same V shows stability because

$$\dot{V} = -x^4$$

However, this "lazy" control design gives very slow convergence near $x = 0$. It is better to add linear damping as in the previous control law, which gives better convergence and the control effort is reasonable.

What does Sontag's universal formula give?

$$a(x) = -x^4, \qquad b(x) = x$$

so we get

$$u = -\frac{-x^4 + \sqrt{x^8 + x^4}}{x^2}x = -\frac{-x^4 + x^2\sqrt{x^4+1}}{x} = x^3 - x\sqrt{x^4+1}$$

(initially we should have defined u separately to be 0 when $x = 0$ but the final formula captures this). The closed-loop system is

$$\dot{x} = -x\sqrt{x^4+1}$$

This control law has a slightly more complicated expression than the previous ones, but it has the following interesting properties. First, we have $u \to 0$ as $|x| \to \infty$, i.e., for large x we do nothing and let the $-x^3$ term do all the work. On the other hand, for small x we have $\dot{x} \approx -x$, which ensures nice convergence to 0. So, this control is a good compromise between the previous designs! □

4.2 Back to the adaptive control example

Back to Example 1: consider the plant and the tuning law but open-loop (control yet to be designed):

$$\dot{x} = \theta x + u$$
$$\dot{\hat{\theta}} = x^2$$

Our Lyapunov function candidate was

$$V(x,\hat{\theta}) = \frac{x^2}{2} + \frac{(\hat{\theta}-\theta)^2}{2}$$

Is this a CLF? Differentiate it:

$$\dot{V} = \theta x^2 + xu + (\hat{\theta}-\theta)x^2 = \hat{\theta}x^2 + xu \tag{38}$$

Does it satisfy the definition of CLF? We have

$$a = \hat{\theta}x^2, \qquad b = x$$

We have to be careful: the answer is no! This is because the state of the system is $(x,\hat{\theta})$, so for V to be a CLF we would need the implication

$$b(x,\hat{\theta}) = 0 \quad \Longrightarrow \quad a(x,\hat{\theta}) < 0$$

to hold for all $(x,\hat{\theta}) \neq (0,0)$. But it only holds when $x \neq 0$. In other words, the subset where it doesn't hold is not just the origin, but the whole line $\{x = 0\}$. This is not surprising because we know it is the line of equilibria.

But since our objective is to stabilize the system only to this line, and not to the origin, the universal formula will still hopefully solve our problem. Actually, the above V is a "weak CLF" in the sense that

$$b(x,\hat{\theta}) = 0 \quad \Longrightarrow \quad a(x,\hat{\theta}) \leq 0$$

everywhere (note the nonstrict inequality). And since $b = x$, we do have control authority to make $\dot{V}$ negative where it matters, i.e., for $x \neq 0$.

Lecture 8 (Mon Sep 28)

$\longrightarrow$ Problem Set 2 is assigned.

The universal formula gives

$$k(x) = \begin{cases} -\dfrac{\hat{\theta}x^2 + \sqrt{\hat{\theta}^2 x^4 + x^4}}{x^2} x, & x \neq 0 \\ 0, & x = 0 \end{cases}$$

which simplifies to

$$k(x) = -\frac{\hat{\theta}x^2 + x^2\sqrt{\hat{\theta}^2 + 1}}{x} = -\left(\hat{\theta} + \sqrt{\hat{\theta}^2 + 1}\right)x$$

Plugging this into the expression (38) for $\dot{V}$, we get

$$\dot{V} = -\sqrt{\hat{\theta}^2 + 1}\, x^2$$

Now Theorem 2 implies $x \to 0$ as before.

The important thing was that a and b did not depend on the unknown parameter θ, so the above control is implementable. Contrast this with

$$V = \frac{x^2}{2}$$

which gives

$$\dot{V} = \theta x^2 + xu$$

and we can't do anything unless we know θ.

We see that the CLF concept does provide useful insight into the adaptive example, but it needs to be tweaked. First, we only have a weak CLF. Second, we have to make sure that $\dot{V}$ does not involve θ. This kind of situation is typical in adaptive control.

We'll see Lyapunov-based design again in the next two sections—backstepping and parameter estimation. The universal formula itself is not used very often in the actual control design, because usually simpler control laws can be found on a case by case basis (as in the above examples). But it is nice to know that a universal formula exists.

5 Backstepping

Reference: [KKK book, Chapters 2 and 3]

We now discuss one of very few available tools for systematically generating Lyapunov functions for certain classes of systems. This is also a nonlinear design technique, and it'll allow us to venture into adaptive control of *nonlinear* plants. Up to now, the only nonlinear control design tool we've discussed is Lyapunov-based design (and Sontag's universal formula), and backstepping is in a way a continuation of that. Plants in all adaptive control scenarios have been linear so far.

In the non-adaptive context, the idea of backstepping first appeared in the Russian literature: paper by Meilakhs, 1978. Independently and around the same time, it was investigated in the adaptive control context (MRAC problem) by Feuer and Morse, and then by Morse and the authors of the KKK book (who subsequently coined the term "backstepping" as a general technique for not necessarily adaptive control).

5.1 Integrator backstepping

Start with the affine control system

$$\dot{x} = f(x) + G(x)u, \qquad x \in \mathbb{R}^n, \ u \in \mathbb{R}^m \tag{39}$$

and assume that we have a CLF $V_0(x)$ and a stabilizing control law

$$u = k_0(x), \qquad k_0(0) = 0$$

for which

$$\frac{\partial V_0}{\partial x} \cdot f(x) + \frac{\partial V_0}{\partial x} \cdot G(x)k_0(x) \le -W(x) < 0 \qquad \forall\, x \neq 0$$

Assume that all data (V_0, f, G, k_0) are smooth.

Now, suppose that our system is augmented with an integrator, increasing relative degree (think of x as output) from 1 to 2:

$$\begin{aligned} \dot{x} &= f(x) + G(x)\xi \\ \dot{\xi} &= u \end{aligned}$$

(Note that ξ is a vector in $\mathbb{R}^m$, so this is an m-dimensional integrator.) We want to find a CLF and a stabilizing feedback law for this new system.

We can view $\xi = k_0(x)$ as a "virtual" control law, which is no longer implementable (because ξ is a state and not a control). Motivated by this, define an "augmented" candidate CLF

$$V_1(x, \xi) := V_0(x) + \frac{1}{2}|\xi - k_0(x)|^2 \tag{40}$$

Its derivative along the (x, ξ)-system is (k_0' stands for the Jacobian matrix $\frac{\partial k_0}{\partial x}$)

$$\begin{aligned}
\dot{V}_1 &= \frac{\partial V_0}{\partial x} \cdot f + \frac{\partial V_0}{\partial x} \cdot G\xi + (\xi - k_0)^T (u - k_0' f - k_0' G\xi) \\
&= \frac{\partial V_0}{\partial x} \cdot f + \frac{\partial V_0}{\partial x} \cdot Gk_0 + \frac{\partial V_0}{\partial x} \cdot G(\xi - k_0) + (\xi - k_0)^T (u - k_0' f - k_0' G\xi) \\
&= \underbrace{\frac{\partial V_0}{\partial x} \cdot f + \frac{\partial V_0}{\partial x} \cdot Gk_0}_{\text{“old” } \dot{V}_0, \text{ for (39)}} + (\xi - k_0)^T \Big(u - k_0' f - k_0' G\xi + G^T \frac{\partial V_0}{\partial x} \Big) \\
&\leq -W(x) + (\xi - k_0)^T \Big(u - k_0' f - k_0' G\xi + G^T \frac{\partial V_0}{\partial x} \Big)
\end{aligned}$$

Claim: V_1 is a CLF. (See the characterization (36) of CLF for affine systems.)

Indeed, the term multiplying u is $(\xi - k_0(x))^T$. Suppose that it is 0. Then what remains is $\leq -W(x)$, hence it is negative unless $x = 0$ (since W is positive definite). If $x = 0$, then $k_0(x) = k_0(0) = 0$ and so to have $\xi - k_0(x) = 0$ we must have $\xi = 0$. We showed that away from $(x, \xi) = (0, 0)$, we can make $\dot{V}_1 < 0$ by a proper choice of u, which proves the claim.

It is also not hard to see how a stabilizing feedback law can be designed. For example, we can simply cancel the terms in large parentheses and add a negative square to $\dot{V}_1$:

$$u = k_1(x, \xi) := -(\xi - k_0) + k_0' f + k_0' G\xi - G^T \frac{\partial V_0}{\partial x}$$

gives

$$\dot{V}_1 \leq -W(x) - |\xi - k_0|^2 < 0 \qquad \forall\, (x, \xi) \neq (0, 0)$$

where the last inequality is proved as in the proof of the previous Claim.

Note that since k_1 involves the derivative of the nominal control law k_0, we lose one degree of smoothness in the control law.

$\longrightarrow$ The key idea of backstepping is not the actual formula for the control law, but the procedure of constructing the augmented Lyapunov function V_1 as in (40). We usually have some flexibility in the choice of the control law, which is common in Lyapunov-based design as we already discussed before. The next example illustrates the procedure and this last point.

Example 4 Consider the 2-D system

$$\begin{aligned}
\dot{x} &= -x^3 + \xi \\
\dot{\xi} &= u
\end{aligned} \tag{41}$$

Rather than just applying the general formula derived above, let's follow the procedure to see better how it works. We first need a CLF and a stabilizing control law $u = k_0(x)$ for the scalar system

$$\dot{x} = -x^3 + u$$

We already considered this system in Example 3 (page 36), where we had the CLF

$$V_0(x) = \frac{x^2}{2}$$

In fact, when the x-system is scalar, a CLF (if one exists) can always be this one. One choice of the control law was

$$k_0(x) = -x$$

which gives

$$\dot{V}_0 = -x^4 - x^2$$

Next, consider

$$V_1(x,\xi) = V_0(x) + \frac{1}{2}(\xi - k_0(x))^2 = \frac{x^2}{2} + \frac{1}{2}(\xi + x)^2$$

We compute its derivative along the (x,ξ)-system:

$$\begin{aligned}\dot{V}_1 &= -x^4 + x\xi + (\xi + x)(u - x^3 + \xi)\\ &= -x^4 - x^2 + x(\xi + x) + (\xi + x)(u - x^3 + \xi)\\ &= -x^4 - x^2 + (\xi + x)(u + x - x^3 + \xi)\end{aligned}$$

hence we can define

$$u = k_1(x,\xi) := -2(\xi + x) + x^3$$

to get

$$\dot{V}_1 = -x^4 - x^2 - (\xi + x)^2 < 0 \qquad \forall\,(x,\xi) \neq (0,0)$$

and we are done.

The above feedback law cancels the x^3 term in $\dot{V}_1$, which is not really necessary. With the CLF V_1 in hand, we have flexibility in selecting another feedback as long as it gives $\dot{V}_1 < 0$. An example of another choice of feedback is

$$u = -2(\xi + x) - x^2(\xi + x)$$

This is still stabilizing because it gives

$$\begin{aligned}\dot{V}_1 &= -x^4 - x^2 - (\xi + x)^2 - x^2(\xi + x)^2 - x^3(\xi + x)\\ &= -\frac{3}{4}x^4 - x^2 - (\xi + x)^2 \underbrace{-x^2(\xi + x)^2 - x^3(\xi + x) - \frac{1}{4}x^4}_{=-\left(x(\xi+x)+\frac{1}{2}x^2\right)^2} < 0 \qquad \forall\,(x,\xi) \neq (0,0)\end{aligned}$$

This control law perhaps makes a little bit more sense because it is close to 0 when $\xi \approx -x$, i.e., when the behavior of ξ is consistent with $k_0(x)$. $\square$

We presented only one step of integrator backstepping. However, it is a recursive procedure which can be applied in the obvious way to a chain of k integrators:

$$\begin{aligned}\dot{x} &= f(x) + G(x)\xi_1 \\ \dot{\xi}_1 &= \xi_2 \\ &\vdots \\ \dot{\xi}_{k-1} &= \xi_k \\ \dot{\xi}_k &= u\end{aligned}$$

We assume that a CLF V_0 for the x-system is given (if x is scalar or has a low dimension, we can hope to find one easily by hand.) Then, we generate a sequence of CLFs $V_1, \dots, V_k$, and the last one is a CLF for the entire system. As before, we lose one degree of smoothness of the feedback law at each step, so we need to make sure that k_0 is at least $\mathcal{C}^k$.

$\longrightarrow$ Wednesday lecture canceled due to Allerton Conference.

Lecture 9 (Mon Oct 5)

It is useful to compare backstepping with feedback linearization on the above example. The system (41) is feedback linearizable. To better explain what we mean by feedback linearization (which was only briefly mentioned in an earlier example illustrating Sontag's formula), define

$$z := -x^3 + \xi$$

and compute its derivative:

$$\dot{z} = -3x^2(-x^3 + \xi) + u = 3x^5 - 3x^2\xi + u$$

If we apply the feedback law

$$u = -3x^5 + 3x^2\xi - x - z$$

then, if we convert everything to the (x, z)-coordinates, we get the closed-loop system

$$\begin{aligned} \dot{x} &= z \\ \dot{z} &= -x - z \end{aligned}$$

which is linear and asymptotically stable. (We can check that the Jacobian of the map from (x, ξ) to (x, z) is nonsingular and hence the coordinate transformation is well defined.)

This particular feedback linearizing controller involves terms of higher degrees than the backstepping one, because it cancels all nonlinearities while the backstepping controller tries to preserve "friendly" nonlinearities. On the other hand, backstepping requires more structure from the system: all "virtual controls" ξ_i must enter affinely on the right-hand sides. Feedback linearization doesn't require this; for example, we can apply it to

$$\begin{aligned} \dot{x} &= f(x, \xi) \\ \dot{\xi} &= u \end{aligned}$$

as long as $\frac{\partial f}{\partial \xi} \neq 0$. For backstepping, the first equation must be of the form

$$\dot{x} = f(x) + g(x)\xi$$

Note, however, that we don't need to assume $g(x) \neq 0$ to use the Lyapunov-based design.

In other words, the two techniques—backstepping and feedback linearization—are complementary as they apply to different classes of systems (although in both cases, relative degree must equal state space dimension).

5.2 Adaptive integrator backstepping

In our old Example 1 we had: the plant

$$\dot{x} = \theta x + u$$

the controller

$$u = -(\hat{\theta} + 1)x =: k_0(x, \hat{\theta})$$

the tuning law

$$\dot{\hat{\theta}} = x^2 =: \tau_0(x)$$

and the Lyapunov function

$$V_0(x, \hat{\theta}) := \frac{1}{2}\big(x^2 + (\hat{\theta} - \theta)^2\big)$$

For convenience let us introduce the parameter error

$$\tilde{\theta} := \hat{\theta} - \theta$$

then we get

$$V_0(x, \hat{\theta}) := \frac{1}{2}\big(x^2 + \tilde{\theta}^2\big)$$

and

$$\dot{V}_0 = \frac{\partial V_0}{\partial x}(\theta x + k_0) + \frac{\partial V_0}{\partial \hat{\theta}}\tau_0(x) = x(\theta x - (\hat{\theta} + 1)x) + \tilde{\theta}x^2 = -x^2$$

Let us now add an integrator:

$$\begin{aligned} \dot{x} &= \theta x + \xi \\ \dot{\xi} &= u \end{aligned}$$

One complication is that the above "virtual" control law is dynamic. However, we can still apply the same idea and consider the augmented candidate CLF

$$V_1(x, \hat{\theta}, \xi) := \frac{1}{2}\big(x^2 + \tilde{\theta}^2 + (\xi - k_0(x, \hat{\theta}))^2\big) = \frac{1}{2}\big(x^2 + \tilde{\theta}^2 + (\xi + (\hat{\theta} + 1)x)^2\big)$$

Let's write its derivative, keeping $\dot{\hat{\theta}}$ open for now since we're not yet sure if the same tuning law will work.

$$\begin{aligned} \dot{V}_1 &= x(\theta x + \xi) + \tilde{\theta}\dot{\hat{\theta}} + (\xi + (\hat{\theta} + 1)x)\big(u + \dot{\hat{\theta}}x + (\hat{\theta} + 1)(\theta x + \xi)\big) \\ &= \underbrace{x(\theta x - (\hat{\theta} + 1)x) + \tilde{\theta}\dot{\hat{\theta}}}_{=-x^2 + \tilde{\theta}(\dot{\hat{\theta}} - x^2)} + (\xi + (\hat{\theta} + 1)x)\big(u + x + \dot{\hat{\theta}}x + (\hat{\theta} + 1)(\theta x + \xi)\big) \end{aligned}$$

Difficulty: if we define $\dot{\hat{\theta}} = x^2$ as before, then, to get $\dot{V}_1 < 0$, we need to define u to cancel the terms in the last parentheses and add a damping term $-(\xi + (\hat{\theta}+1)x)$. This is what we did in the non-adaptive case. But the terms in the parentheses depend on the unknown parameter θ!

Solution: Replace θ by $\hat{\theta}$ in the last parentheses, and the carry the difference—which depends on $\tilde{\theta}$—outside the parentheses and combine it with the $\tilde{\theta}\dot{\hat{\theta}}$ term.

$$\dot{V}_1 = -x^2 + \tilde{\theta}\big(\dot{\hat{\theta}} - x^2 - (\xi + (\hat{\theta}+1)x)(\hat{\theta}+1)x\big) + (\xi + (\hat{\theta}+1)x)\big(u + x + \dot{\hat{\theta}}x + (\hat{\theta}+1)(\hat{\theta}x + \xi)\big)$$

Now the choice of the tuning law and the control law is clear: first set

$$\dot{\hat{\theta}} = x^2 + (\xi + (\hat{\theta}+1)x)(\hat{\theta}+1)x$$

which cancels the $\tilde{\theta}$ term. Note that this takes the form

$$\dot{\hat{\theta}} = \tau_0 + \tau_1 =: \tau$$

where $\tau_0 = x^2$ is what we had for the scalar plant and τ_1 is a new term. Next, set

$$u = -x - \tau x - (\hat{\theta}+1)(\hat{\theta}x + \xi) - (\xi + (\hat{\theta}+1)x) =: k_1(x, \hat{\theta}, \xi)$$

This gives

$$\dot{V}_1 = -x^2 - (\xi + (\hat{\theta}+1)x)^2 =: -W_1(x, \hat{\theta}, \xi) < 0 \qquad \text{when } (x, \xi) \neq (0,0)$$

and we are in good shape because this implies that $x, \hat{\theta}, \xi$ are bounded and $x, \xi \to 0$ (Theorem 2).

$\longrightarrow$ Note that the above controller is not really a certainty equivalence controller any more: it incorporates explicitly a correction term (τx) coming from the tuning law. (Compare this with the discussion on page 8.)

We see—and this is not surprising—that adaptive backstepping proceeds along the same lines as non-adaptive backstepping but it is more challenging.

The next step would be to try to stabilize

$$\begin{aligned} \dot{x} &= \theta x + \xi_1 \\ \dot{\xi}_1 &= \xi_2 \\ \dot{\xi}_2 &= u \end{aligned}$$

The above example suggests considering the candidate CLF

$$V_2(x, \hat{\theta}, \xi_1, \xi_2) := V_1(x, \hat{\theta}, \xi_1) + \frac{1}{2}(\xi_2 - k_1(x, \hat{\theta}, \xi))^2$$

and a tuning law of the form

$$\dot{\hat{\theta}} = \tau_0 + \tau_1 + \tau_2$$

The functions $\tau_0, \tau_1, \tau_2, \ldots$ are called *tuning functions*. For a general procedure of designing them, see [KKK book, Chapter 4].

6 Parameter estimation

There are no nonlinear system theory concepts here, but instead some ideas from optimization algorithms and the important concept of a persistently exciting signal.

General points:

- Unknown parameters are often difficult to estimate off-line, hence the need for *on-line estimation*. Examples: camera calibration; calibration of mechanical system models (such as helicopter in flight).
- We will see that convergence of parameter estimates to their true values requires some special properties of the input signal u.
- Sometimes parameter convergence is not crucial, and u is chosen to meet behavior specs of the system. (For example, to fly a helicopter, it is not necessary to know all its model parameters exactly.) To analyze stability of the resulting adaptive control system, other properties of the parameter estimation scheme (in particular, slow adaptation speed) will be important.

Choice of control—later. Here—on-line parameter estimation. Reference: [Ioannou-Sun, Chap. 4,5]. Most of this material is also in [Ioannou-Fidan].

A word on terminology: "Parameter estimation" is some procedure for generating parameter estimates, whether or not they converge to the true values of the parameters. "Parameter identification" is a stronger term, which assumes that parameter estimates do converge to the true values. (This is the difference between Chapters 4 and 5 in [Ioannou-Sun]; in [Ioannou-Fidan] this is combined in one chapter.) For control purposes, we will always use some form of parameter estimation, but will not necessarily require parameter identification. Of course, some other nice properties of the parameter estimation scheme will be needed in such cases. By the end of the course, we will understand why adaptive control works even without parameter identification (and will thus justify control design based on the certainty equivalence principle).

Lecture 10 (Wed Oct 7)

6.1 Gradient method

Parameter estimation algorithms usually aim at minimizing some cost function that reflects the quality of estimation. Let us recall some basic facts about minimizing a given function $f : \mathbb{R}^n \to \mathbb{R}$. We say that x^* is a (global) minimum of f if

$$f(x^*) \le f(x) \qquad \forall\, x$$

If $f \in \mathcal{C}^1$ then its *gradient* is the vector

$$\nabla f(x) := \left(\frac{\partial f}{\partial x_1}(x), \dots, \frac{\partial f}{\partial x_n}(x)\right)^T$$

A necessary condition for x^* to be a global minimum is

$$\nabla f(x^*) = 0 \tag{42}$$

This condition is in general not sufficient, because it also holds at *local* minima, as well as at maxima and saddle points. However, if f is a *convex* function, then the above 1st-order condition is sufficient (for convex functions, a stationary point is automatically a global minimum).

The most basic algorithm[3] for computing the minimum is the *gradient method* (or the method of *steepest descent*). Its continuous version is given by the differential equation

$$\dot{x} = -\nabla f(x) \tag{43}$$

(with some initial condition x_0). More generally, one considers

$$\dot{x} = -\Gamma \nabla f(x) \tag{44}$$

where

$$\Gamma = \Gamma^T > 0$$

is a *scaling matrix*. This is exactly what we get from the original (unscaled) gradient method by changing coordinates in $\mathbb{R}^n$ according to

$$x = \Gamma_1 \bar{x}, \qquad \Gamma_1 \Gamma_1^T = \Gamma \tag{45}$$

(See, e.g., [Ioannou-Sun, Section B.2, p. 785].)

The idea behind the gradient method is that f decreases along trajectories of (44), and eventually we approach an equilibrium where (42) holds. Convergence of the gradient method is a

[3]There are of course others, such as Newton's method.

standard topic in optimization textbooks, although results are more often stated for discrete iterations (with an appropriate choice of stepsize). For (43), we can consider the candidate Lyapunov function

$$V(x) := f(x) - f(x^*)$$

where x^* is a global minimum (assumed to exist). Its derivative is

$$\dot{V}(x) = -|\nabla f(x)|^2$$

hence we have

$$\nabla f(x) \to 0$$

along any bounded solution (cf. Theorem 2). For (44), the same V gives

$$\dot{V}(x) = -\nabla f(x)\Gamma\nabla f(x) \tag{46}$$

and the same conclusion holds since Γ is positive definite.

In view of earlier remarks, (46) does not necessarily imply that $x(t)$ converges to the global minimum x^*. But if f is a convex function, then we do in fact have convergence to a global minimum.

If we know that f is convex, then we can use another Lyapunov function:

$$V(x) := |x - x^*|^2$$

which can be shown to decrease along solutions of (43). (Reason: if $f(x^*) < f(x)$ and f is convex, then f decreases as we start moving along the line from x to x^*. Hence, the inner product with the gradient is negative.) To handle the scaled gradient law (44), we need to consider

$$V(x) := (x - x^*)^T\Gamma^{-1}(x - x^*) \tag{47}$$

which is the same as $|\bar{x} - \bar{x}^*|^2$ for $\bar{x}$ given by (45).

$\longrightarrow$ We will see Lyapunov functions of this form several times below when analyzing convergence of parameter estimation schemes based on gradient laws. However, since the gradient law will be coupled with plant dynamics, we will need to do the proof from scratch.

6.2 Parameter estimation: stable case

We first discuss stable plants with bounded inputs. Afterwards, we'll discuss how to lift these assumptions.

Example 5 Let us start with the static scalar example

$$y(t) = \theta u(t)$$

where $\theta \in \mathbb{R}$ is the unknown parameter and $u(\cdot)$, $y(\cdot)$ are known continuous signals. For now we assume that u is *bounded* (hence so is y). We'll sometimes write this as $u \in L_\infty$, $y \in L_\infty$, etc.

Problem: estimate θ.

A naive approach is just to use

$$\theta = \frac{y}{u}$$

However, this has problems. First, it's ill-defined if $u = 0$. Second, this is very sensitive to noise. But most importantly, we want a method that will work for the dynamic case (when u and y are the input and output of a dynamical system containing uncertain parameter θ).

So, instead we want to design a dynamical system which will generate a time-varying estimate $\hat{\theta}(t)$ of θ. This will lead also to

$$\hat{y}(t) := \hat{\theta}(t)u(t)$$

Note: $\hat{y}$ is not really an estimate of y, since we can measure y directly. But $\hat{y}$ will provide us feedback on the quality of the estimate $\hat{\theta}$. To this end, we define the output estimation (or output prediction) error

$$e(t) := \hat{y}(t) - y(t) = (\hat{\theta}(t) - \theta)u(t)$$

(in Ioannou-Sun, e is defined with the opposite sign). Based on this error e, we define the *cost function*

$$J(\hat{\theta}(t), t) := \frac{e^2(t)}{2} = \frac{(\hat{\theta} - \theta)^2 u^2}{2}$$

For simplicity, we will omit the argument t and just write $J(\hat{\theta})$.

Idea: update $\hat{\theta}$ so as to minimize $J(\hat{\theta})$.

The motivation is that $J \geq 0$ always and $J = 0$ when $\hat{\theta} = \theta$. Note, however, that it's possible to have $J = 0$ when $\hat{\theta} \neq \theta$ if $u = 0$. To avoid this, we make the (temporary) assumption that

$$u^2 \geq c > 0 \tag{48}$$

We want to use the gradient method.

$$\nabla J(\hat{\theta}) = (\hat{\theta} - \theta)u^2 = eu$$

Thus the gradient law is

$$\dot{\hat{\theta}} = -\gamma e u \tag{49}$$

with some initial condition $\hat{\theta}_0$, where $\gamma > 0$ is the scaling factor (or "adaptive gain").

Note that J is convex as a function of $\hat{\theta}$, thus there's hope for convergence. But J also depends on $u(t)$, so this is not quite as easy as the standard case discussed earlier. Candidate Lyapunov function

$$V(\hat{\theta}) := \frac{(\hat{\theta} - \theta)^2}{2} \tag{50}$$

gives

$$\dot{V} = (\hat{\theta} - \theta)\dot{\hat{\theta}} = -(\hat{\theta} - \theta)\gamma e u = -\gamma(\hat{\theta} - \theta)^2 u^2 \leq -\gamma(\hat{\theta} - \theta)^2 c < 0 \qquad \forall\, \hat{\theta} \neq \theta$$

where c comes from (48). This implies that $\hat{\theta} = \theta$ is an asymptotically stable equilibrium. (In fact, it is exponentially stable since $\dot{V} \leq -2\gamma c V$ and V is quadratic.)

Now, suppose we don't want to assume (48), but allow instead any bounded u *with bounded derivative*. The previous formulas for $\dot{V}$ can be rewritten as

$$\dot{V} = -\gamma e^2$$

We can now use the same trick as in the proof of Theorem 2: integrate this and reduce to Barbalat's lemma.

$$V(t) - V(0) = -\gamma \int_0^t e^2(s) ds \quad \Longrightarrow \quad \gamma \int_0^t e^2(s) ds = V(0) - V(t) \leq V(0) \qquad \forall\, t$$

which implies that $\int_0^\infty e^2$ is finite, i.e., $e \in L_2$.

To apply Barbalat, we need to show that e, $\dot{e}$ are bounded.

Since V is nonincreasing, from the definition (50) of V it is clear that $\hat{\theta}$ is bounded.

Thus $e = (\hat{\theta} - \theta)u$ is also bounded.

In view of (49), this shows that $\dot{\hat{\theta}}$ is bounded, and belongs to L_2 as well[4].

We then have that $\dot{e} = \dot{\hat{\theta}} u + (\hat{\theta} - \theta)\dot{u}$ is also bounded.

By Barbalat's lemma, $e \to 0$, i.e., $\hat{y} \to y$, and $\dot{\hat{\theta}} \to 0$.

Note: we didn't use Theorem 2 directly because of the presence of the external input u, but we essentially repeated the steps of its proof.

What does the above analysis imply about convergence of $\hat{\theta}$?

Since V is nonincreasing and bounded from below by 0, it has a limit as $t \to \infty$. Hence, $\hat{\theta}$ must converge to a constant. (Note that this doesn't follow just from the fact that $\dot{\hat{\theta}} \to 0$.)

Can we conclude from the above analysis that $\hat{\theta} \to \theta$?

Not necessarily. Define the parameter estimation error

$$\tilde{\theta} := \hat{\theta} - \theta$$

It satisfies the DE

$$\dot{\tilde{\theta}} = \dot{\hat{\theta}} = -\gamma e u = -\gamma u^2 \tilde{\theta}$$

[4]Such "slow adaptation" properties will be useful later when we discuss stability of slowly time-varying systems.

which we can solve to get

$$\tilde{\theta}(t) = e^{-\gamma \int_0^t u^2(s)ds}\tilde{\theta}(0)$$

It is now clear that convergence of $\tilde{\theta}$ to 0 is affected by the choice of u. For example, if u satisfies (48), then clearly $\tilde{\theta}$ converges exponentially fast (at the rate at least γc), as we already proved earlier. On the other hand, if $u \in L_2$, then the integral in the exponent is bounded and $\tilde{\theta}$ does not converge.

Lecture 11 (Mon Oct 12)

Let us try to make precise the convergence property we want from $\tilde{\theta}$ and the corresponding condition that should be required from u. We say that $\tilde{\theta}$ is *uniformly exponentially convergent (UEC)* if for some $c, \lambda > 0$ we have

$$\tilde{\theta}(t) \leq ce^{-\lambda(t-t_0)}\tilde{\theta}(t_0)$$

for all $t \geq t_0 \geq 0$. We say that u is *persistently exciting (PE)* if for some $\alpha_0, T_0 > 0$ we have

$$\int_t^{t+T_0} u^2(s)ds \geq \alpha_0 T_0 \qquad \forall\, t \tag{51}$$

This is an important concept (we'll extend it later to vector signals). Constant signals, or more generally signals satisfying (48) are PE, while L_2 signals are not. The constant α_0 is called the *level of excitation.*

Lemma 4 *For the above gradient law, $\tilde{\theta}$ is UEC if and only if u is PE.*

Proof—next HW.

□

Example 6 Consider now the one-dimensional plant

$$\dot{x} = -ax + bu$$

where we assume that $a > 0$ and u is bounded (hence x is also bounded). We want to estimate the unknown parameters a and b.

Estimator[5]*:*

$$\dot{\hat{x}} = -a_m(\hat{x} - x) - \hat{a}x + \hat{b}u$$

where $a_m > 0$ is a design constant which determines the damping rate of the estimator. To see this, define the estimation errors

$$e := \hat{x} - x, \qquad \tilde{a} := \hat{a} - a, \qquad \tilde{b} := \hat{b} - b$$

(Again, note that calling e an "estimation error" is not really accurate because both $\hat{x}$ and x are measured signals; sometimes the term "state prediction error" is used instead.) Then e satisfies the DE

$$\dot{e} = -a_m e - \tilde{a}x + \tilde{b}u$$

(this DE is not actually implemented, but only used for analysis purposes).

---

[5]This is not the only possible estimator; see [Ioannou-Sun, Sect. 4.2.2] for another estimator and comparison.

Observe that the above equation is stable with respect to x, u and has the autonomous contraction rate a_m. In particular, if $\tilde{a}, \tilde{b}$ are 0 or converge to 0, then $e \to 0$ (recall that x, u are bounded). However, the converse is not true: $e \to 0$ does not necessarily imply $\tilde{a}, \tilde{b} \to 0$ unless the signals $x, \hat{x}, u$ are PE in some sense (we will see this later).

Update laws for the estimates $\hat{a}, \hat{b}$ will be driven by e, and will take the form

$$\dot{\hat{a}} = f_1(e, \hat{x}, x, u), \qquad \dot{\hat{b}} = f_2(e, \hat{x}, x, u)$$

Goal: make $\tilde{a}, \tilde{b}, e \to 0$ along solutions of the resulting 3-D system.

Lyapunov-based design: Consider the candidate Lyapunov function

$$V(e, \tilde{a}, \tilde{b}) := \frac{1}{2}(e^2 + \tilde{a}^2 + \tilde{b}^2) \tag{52}$$

(has all desired properties: C^1, positive definite, radially unbounded).

$$\dot{V} = e\dot{e} + \tilde{a}\dot{\hat{a}} + \tilde{b}\dot{\hat{b}} = e(-a_m e - \tilde{a}x + \tilde{b}u) + \tilde{a}f_1 + \tilde{b}f_2 = -a_m e^2 \underbrace{-\tilde{a}xe + \tilde{b}ue + \tilde{a}f_1 + \tilde{b}f_2}_{\text{not helpful, want to cancel this}}$$

So, a natural choice is

$$f_1 := xe, \qquad f_2 := -ue \qquad \Longrightarrow \qquad \dot{V} = -a_m e^2 \le 0$$

What conclusions can we draw from this?

V is bounded $\Longrightarrow$ $e, \hat{a}, \hat{b}$ are bounded.

$u, x, \hat{x} (= x + e)$ are bounded $\Longrightarrow$ all signals are bounded.

Integrating $\dot{V} = -a_m e^2$ as before we get $e \in L_2$.

$\dot{e}$ is bounded $\Longrightarrow$ by Barbalat's lemma, $e \to 0$. (Can also appeal directly to Theorem 2.)

Thus $\dot{\hat{a}}, \dot{\hat{b}} \to 0$, and they are also in L_2.

Since V is nonincreasing and bounded from below by 0, it has a limit as $t \to \infty$. And, since we know that $e \to 0$, we get that $\tilde{a}^2 + \tilde{b}^2$ also has a limit. However, this doesn't mean that $\tilde{a}$ and $\tilde{b}$ individually converge to any values, let alone to 0.

So, the above parameter estimation scheme has a number of useful properties but it does not guarantee convergence of parameter estimates $\hat{a}, \hat{b}$ to their true values a, b. This is actually not surprising. For example, if $u \equiv 0$ and $x_0 = 0$ then x stays at 0 and we cannot learn a, b no matter what we do. In other words, some PE-like assumptions will have to be imposed on the input u to guarantee convergence of parameter estimates (we already saw this in the previous example). □

$\longrightarrow$ These examples are just to build intuition. We'll make things more rigorous and more general a little later.

We now consider the general vector case, but still assuming stability. Namely, the plant is

$$\dot{x} = Ax + Bu$$

where $x \in \mathbb{R}^n$, $u \in \mathbb{R}^m$ is bounded, and A, B are unknown matrices with A Hurwitz (hence we know that x is bounded as well). We are assuming here that x is available for measurement (i.e., $y = x$).

Estimator is an extension of the one from the previous scalar example:

$$\dot{\hat{x}} = A_m(\hat{x} - x) + \hat{A}x + \hat{B}u$$

where A_m is a Hurwitz matrix (chosen by the designer) and $\hat{A}, \hat{B}$ are estimates of A, B (to be generated). Defining as before

$$e := \hat{x} - x, \qquad \tilde{A} := \hat{A} - A, \qquad \tilde{B} := \hat{B} - B$$

we have

$$\dot{e} = A_m e + \tilde{A}x + \tilde{B}u$$

Update laws:

$$\dot{\hat{A}} = F_1(e, \hat{x}, x, u), \qquad \dot{\hat{B}} = F_2(e, \hat{x}, x, u)$$

Candidate Lyapunov function:

$$V(e, \tilde{A}, \tilde{B}) := e^T P e + \mathrm{tr}(\tilde{A}^T \tilde{A}) + \mathrm{tr}(\tilde{B}^T \tilde{B})$$

where P comes from the Lyapunov equation

$$A_m^T P + P A_m = -I$$

and tr stands for trace of a matrix. Note that $\mathrm{tr}(\tilde{A}^T \tilde{A})$ is nothing but the sum of squares of all elements of $\tilde{A}$.

$\longrightarrow$ In class we skipped the calculation that follows.

We have

$$\begin{aligned}
\dot{V} &= (e^T A_m^T + x^T \tilde{A}^T + u^T \tilde{B}^T) P e + e^T P (A_m e + \tilde{A}x + \tilde{B}u) \\
&\quad + \mathrm{tr}(\dot{\tilde{A}}^T \tilde{A} + \tilde{A}^T \dot{\tilde{A}}) + \mathrm{tr}(\dot{\tilde{B}}^T \tilde{B} + \tilde{B}^T \dot{\tilde{B}}) \\
&= e^T (A_m^T P + P A_m) e + x^T \tilde{A}^T P e + u^T \tilde{B}^T P e + e^T P \tilde{A} x + e^T P \tilde{B} u \\
&\quad + \mathrm{tr}(F_1^T \tilde{A} + \tilde{A}^T F_1) + \mathrm{tr}(F_2^T \tilde{B} + \tilde{B}^T F_2) \\
&= -e^T e + 2x^T \tilde{A}^T P e + 2u^T \tilde{B}^T P e + 2\mathrm{tr}(\tilde{A}^T F_1) + 2\mathrm{tr}(\tilde{B}^T F_2)
\end{aligned}$$

Useful property of trace: for two vectors $k, l \in \mathbb{R}^n$ we have $\mathrm{tr}(kl^T) = k^T l = l^T k$ (quick exercise: prove this). Hence

$$\underbrace{x^T}_{l^T} \underbrace{\tilde{A}^T P e}_{k} = \mathrm{tr}(\tilde{A}^T P e x^T)$$

and similarly

$$u^T \tilde{B}^T Pe = \mathrm{tr}(\tilde{B}^T Peu^T)$$

So we get

$$\dot{V} = -e^T e + 2\mathrm{tr}(\tilde{A}^T Pex^T + \tilde{A}^T F_1) + 2\mathrm{tr}(\tilde{B}^T Peu^T + \tilde{B}^T F_2)$$

This makes the choice of F_1, F_2 obvious:

$$F_1 := -Pex^T, \qquad F_2 := -Peu^T \qquad \Longrightarrow \qquad \dot{V} = -|e|^2$$

Boundedness of all signals follows right away.

As before, we have $e \in L_2$, $\dot{e} \in L_\infty \implies e \to 0$.

$\dot{\hat{A}}, \dot{\hat{B}} \to 0$, and they are also in L_2.

To have $\hat{A} \to A$, $\hat{B} \to B$ we will need to place extra assumptions on u to guarantee that it sufficiently excites the plant. Will come back to this later.

6.3 Unstable case: adaptive laws with normalization

Up to now: stable plants, bounded inputs. Not adequate for adaptive control. Want to handle unstable plants, unbounded inputs.

Let us revisit Example 5:

$$y(t) = \theta u(t)$$

$\theta \in \mathbb{R}$ unknown, u, y scalar, *not necessarily bounded* any more.

The basic idea is to *normalize* u and y to get bounded signals. Namely, consider

$$\bar{u} := \frac{u}{m}, \qquad \bar{y} = \frac{y}{m}$$

where m is some signal that guarantees boundedness of $\bar{u}$ (and consequently that of $\bar{y}$). A simple choice is

$$m := \sqrt{1 + u^2}$$

The normalized signals satisfy the same relation:

$$\bar{y} = \theta \bar{u}$$

and we can proceed as before to define the output estimate/predictor

$$\hat{\bar{y}} := \hat{\theta} \bar{u}$$

the output estimation (prediction) error

$$\bar{e} := \hat{\bar{y}} - \bar{y} = \frac{e}{m}$$

(recall that $e = \hat{y} - y$) and the cost

$$J(\hat{\theta}) := \frac{\bar{e}^2}{2} = \frac{(\hat{\theta} - \theta)^2 u^2}{2m^2}$$

We want to use the gradient method to minimize $J(\hat{\theta})$.

$$\nabla J(\hat{\theta}) = \frac{(\hat{\theta} - \theta)u^2}{m^2} = \frac{eu}{m^2}$$

Let us define

$$e_n := \frac{e}{m^2}$$

as this quantity will appear frequently in calculations below. It is called the *normalized* output estimation error. (In [Ioannou-Sun] it is denoted by ϵ.)

Note: at this point we don't yet really need e_n, could also write everything in terms of $\bar{e}, \bar{u}$ (which would look closer to the unnormalized case; exercise: do this) but we'll be relying on this notation later.

The gradient law is

$$\dot{\hat{\theta}} = -\gamma e_n u$$

where $\gamma > 0$. Introduce the parameter estimation error

$$\tilde{\theta} := \hat{\theta} - \theta$$

to have

$$\dot{\tilde{\theta}} = \dot{\hat{\theta}} = -\gamma e_n u = -\gamma \frac{(\hat{\theta} - \theta)u^2}{m^2} = -\gamma \bar{u}^2 \tilde{\theta}$$

Let's try the candidate Lyapunov function

$$V(\tilde{\theta}) := \frac{\tilde{\theta}^2}{2}$$

We have

$$\dot{V} = \tilde{\theta}\dot{\tilde{\theta}} = -\gamma \bar{u}^2 \tilde{\theta}^2 = -\gamma \frac{e^2}{m^2} = -\gamma e_n^2 m^2 \leq 0$$

From this we know how to quickly deduce the following:

$\tilde{\theta} \in L_\infty$ and $e_n m \in L_2$. Since $e_n m = \bar{u}\tilde{\theta}$, it is also in L_∞.

$\dot{\tilde{\theta}} = -\gamma e_n m \bar{u} \in L_2 \cap L_\infty$. I.e., speed of adaptation is bounded in L_2 and L_∞ sense.

$\frac{d}{dt}(e_n m) = \frac{d}{dt}(\tilde{\theta}\bar{u}) = \dot{\tilde{\theta}}\bar{u} + \tilde{\theta}\dot{\bar{u}}$. By Barbalat, $e_n m \to 0$ under the *additional assumption* that $\dot{\bar{u}} \in L_\infty$ (which we also needed in the stable static case, but not in the stable dynamic case). This would in turn imply $\dot{\tilde{\theta}} \to 0$.

We see that basically, normalization lets us recover the main properties of the unnormalized scheme, namely:

- bounded parameter estimates (or, what is the same, bounded parameter estimation errors)
- bounded speed of adaptation in L_2 and L_∞ sense

even though the input is no longer bounded. (But we no longer have convergence of the output estimation error to 0.)

If $m = 1$ (no normalization), then the first item still holds ($\dot{V}$ is still ≤ 0) but the second item doesn't. And we will see later that slow adaptation is important for stability of adaptive control (when estimation is combined with certainty equivalence control).

Of course, without additional assumptions we cannot guarantee that $\tilde{\theta}$ converges to 0 (or to anything else).

Lecture 12 (Wed Oct 14)

In Example 6, we had

$$\dot{x} = -ax + bu$$

but we no longer want to assume $a > 0$ or $u \in L_\infty$. Could try to use normalization, but it's not as straightforward: if we define

$$\bar{u} := \frac{u}{m}, \qquad \bar{x} := \frac{x}{m}$$

where m is chosen so that $\bar{u} \in L_\infty$, then we don't have any simple relationship between $\bar{u}$ and $\bar{x}$. We certainly don't have

$$\dot{\bar{x}} = -a\bar{x} + b\bar{u} \qquad \text{(WRONG)}$$

because m is time-varying (it cannot be constant unless $u \in L_\infty$ in the first place) and so $\dot{m}$ will appear in $\dot{\bar{x}}$.

In [Ioannou-Sun, Sect. 4.3.2] an approach to this problem is proposed which relies on working with a different state estimation error produced by a special system. Rather than trying to fix this example, we now discuss how the general case can be made to resemble the static case of Example 5, so that the previous design can be extended to it.

$\longrightarrow$ Later we will see another approach, which does rely on estimator design (see indirect MRAC). We will return to it again in switching adaptive control. What will happen is that estimator design will be combined with control design, and stabilizing properties of the controller will let us show boundedness of signals, something we cannot do here for arbitrary control inputs.

6.3.1 Linear plant parameterizations (parametric models)

We now want to study a general SISO plant

$$\begin{aligned} \dot{x} &= Ax + bu \\ y &= c^T x \end{aligned}$$

where $x \in \mathbb{R}^n$, the measured signals are $u, y \in \mathbb{R}$, A is an unknown matrix, and b, c are unknown vectors. It is more convenient for us here to represent it in the form

$$y^{(n)} + a_{n-1}y^{(n-1)} + \cdots + a_0 y = b_{n-1}u^{(n-1)} + \cdots + b_0 u \tag{53}$$

The numbers a_i, b_i are the coefficients of the denominator and the numerator of the plant transfer function, respectively:

$$y(s) = \frac{b_{n-1}s^{n-1} + \cdots + b_1 s + b_0}{s^n + a_{n-1}s^{n-1} + \cdots + a_1 s + a_0}u(s) = c^T(Is - A)^{-1}bu(s)$$

We can define

$$\theta := (b_{n-1}, \dots, b_0, a_{n-1}, \dots, a_0)^T, \qquad Y := (u^{(n-1)}, \dots, u, -y^{(n-1)}, \dots, -y)^T$$

and solve (53) for the highest derivative to get

$$y^{(n)} = \theta^T Y \tag{54}$$

This already looks familiar, but the problem is that $y^{(n)}$ and the vector Y cannot be obtained without differentiation. To fix this, filter both sides of (54) with an n-th order stable filter

$$\frac{1}{\Lambda(s)} = \frac{1}{s^n + \lambda_{n-1}s^{n-1} + \dots + \lambda_0}$$

In other words, consider

$$z := \frac{1}{\Lambda(s)} y^{(n)} = \frac{s^n}{\Lambda(s)} y$$

and

$$\phi := \frac{1}{\Lambda(s)} Y = \left(\frac{s^{n-1}}{\Lambda(s)} u, \dots, \frac{1}{\Lambda(s)} u, -\frac{s^{n-1}}{\Lambda(s)} y, \dots, -\frac{1}{\Lambda(s)} y \right)^T \tag{55}$$

The vector ϕ is called the *regressor vector*. It is clear that both the scalar z and the vector ϕ can be generated as outputs of suitable causal linear systems with inputs u and y. (Realizations of proper transfer functions are covered in ECE 515.) They are related by

$$z = \theta^T \phi \tag{56}$$

which is the parameterization we were looking for. It is called a linear parametric model because the unknown parameters enter it linearly (this is unfortunately crucial for most adaptive laws to work[6]).

$\longrightarrow$ The number of parameters in the vector θ is at most $2n$, whereas the number of original parameters (entries of A, b, c was $n^2 + 2n$). This difference is due to the fact that we are only measuring the input/output data, not the state. Different plants with the same transfer function are not distinguished by the parametric model (56).

$\longrightarrow$ The initial condition x_0 is ignored by (56). However, since $\Lambda(s)$ is stable, the contribution of x_0 decays exponentially with time and doesn't affect any of the results below. For details, see [Ioannou-Sun, Sect. 4.3.7].

6.3.2 Gradient method

Estimate/prediction of z:

$$\hat{z} := \hat{\theta}^T \phi$$

[6]This restriction is lifted in switching adaptive control, to be discussed towards the end of the course.

where $\hat{\theta}$ is the estimate of θ as usual. Normalized estimation error:

$$e_n := \frac{\hat{z} - z}{m^2}$$

where

$$m^2 = 1 + n_s^2 \tag{57}$$

and n_s is a normalizing signal such that $\frac{\phi}{m} \in L_\infty$. For example: $n_s = \sqrt{\phi^T \phi}$, or $n_s = \sqrt{\phi^T P \phi}$ for some $P = P^T > 0$. Parameter estimation error:

$$\tilde{\theta} := \hat{\theta} - \theta$$

and we have

$$e_n = \frac{\tilde{\theta}^T \phi}{m^2} \tag{58}$$

Define the cost function

$$J(\hat{\theta}) := \frac{(e_n m)^2}{2} = \frac{(\tilde{\theta}^T \phi)^2}{2m^2} = \frac{((\hat{\theta} - \theta)^T \phi)^2}{2m^2}$$

Intuition: $e_n m = \tilde{\theta}^T \frac{\phi}{m}$, the second term is bounded, so the cost is big when $\tilde{\theta}$ is big.

Also note that J is convex in $\hat{\theta}$.

$$\nabla J(\hat{\theta}) = \frac{\tilde{\theta}^T \phi}{m^2} \phi = e_n \phi$$

Gradient adaptive law:

$$\dot{\hat{\theta}} = -\Gamma e_n \phi$$

where $\Gamma = \Gamma^T > 0$ is a scaling matrix (adaptive gain).

Theorem 5 (Ioannou-Sun, Theorem 4.3.2[7])
(i) $\hat{\theta} \in L_\infty$ *and* $e_n, e_n n_s, \dot{\hat{\theta}} \in L_2 \cap L_\infty$.
(ii) If $\frac{\phi}{m}$ *is PE, then* $\hat{\theta} \to \theta$ *exponentially fast.*

For item (ii), the definition of a persistently exciting (PE) signal needs to be extended to vector-valued signals. This is a direct extension of the earlier scalar definition (51) but instead of squaring the signal we form a *matrix* by multiplying the vector signal on the right by its transpose, and the inequality becomes a matrix one. In other words, $\frac{\phi}{m}$ is PE if for some $\alpha_0, T_0 > 0$ we have

$$\int_t^{t+T_0} \frac{\phi(s)\phi^T(s)}{m^2} ds \geq \alpha_0 T_0 I \qquad \forall\, t$$

[7] The theorem statement there is different, but the same proof works. See also [Ioannou-Fidan, Theorem 3.6.1].

Note that $\phi\phi^T$ is a singular matrix at each time instant (it has rank 1). PE means that over time, it varies in such a way that its integral is uniformly positive definite (ϕ "generates energy in all directions"). When defining PE, one usually requires an upper bound as well as a lower bound:

$$\alpha_1 T_0 I \geq \int_t^{t+T_0} \frac{\phi(s)\phi^T(s)}{m^2} ds \geq \alpha_0 T_0 I$$

but since $\frac{\phi}{m}$ is bounded, the first inequality is automatic.

Proof of Theorem 5.

(i) Take the candidate Lyapunov function (cf. (47))

$$V(\tilde{\theta}) := \frac{1}{2}\tilde{\theta}^T \Gamma^{-1} \tilde{\theta} \tag{59}$$

whose derivative is (recall that Γ is symmetric)

$$\dot{V} = \tilde{\theta}^T \Gamma^{-1} \dot{\tilde{\theta}} = -\tilde{\theta}^T e_n \phi = -e_n^2 m^2 \tag{60}$$

where the last step follows from (58). Now the usual steps:

$\tilde{\theta} \in L_\infty$, hence by (58) $e_n m \in L_\infty$. By (57), $m^2 = 1 + n_s^2$ and so $e_n, e_n n_s \in L_\infty$.

Next, $e_n m \in L_2$. Recalling (57) again, we get $e_n, e_n n_s \in L_2$.

Write

$$\dot{\tilde{\theta}} = -\Gamma e_n m \frac{\phi}{m}$$

and use $e_n m \in L_2 \cap L_\infty$ and $\frac{\phi}{m} \in L_\infty$ to conclude that $\dot{\tilde{\theta}} \in L_2 \cap L_\infty$.

(ii) Using (58), we can write

$$\begin{aligned} \dot{\tilde{\theta}} &= -\Gamma \frac{\phi\phi^T}{m^2}\tilde{\theta} =: A(t)\tilde{\theta} \\ e_n m &= \frac{\phi^T}{m}\tilde{\theta} =: c^T(t)\tilde{\theta} \end{aligned} \tag{61}$$

View this as an LTV system with state $\tilde{\theta}$ and output $e_n m$. The right-hand side of (60) is *minus the output squared*, which is exactly the situation considered in Section 2.2 (the slight difference in notation is because the output in (61) is scalar: c^T here corresponds to C there). By the result stated there, the system (61) will be exponentially stable if we can show that this system is uniformly completely observable (UCO).

To check UCO, we need to analyze the observability Gramian of (61). This looks complicated. However, there is a trick:

Lemma 6 (Ioannou-Sun, Lemma 4.8.1) *A pair $(A(t), C(t))$ is UCO if the pair $(A(t)+L(t)C(t), C(t))$ is UCO for some bounded "output injection" matrix $L(t)$.*

For LTI systems, this is clear from the rank condition for observability. For LTV systems, the proof is harder. See [Ioannou-Sun] for details (the proof is also in the book by Sastry and Bodson). $L(t)$ doesn't actually need to be bounded, but for us this is good enough.

Continuing with the proof of the Theorem: take

$$L(t) := \Gamma \frac{\phi}{m}$$

This gives

$$A(t) + L(t)c^T(t) = 0$$

The observability Gramian of the new pair $\left(0, \frac{\phi^T}{m}\right)$ is

$$M(t_0, t_0 + T) = \int_{t_0}^{t_0+T} \frac{\phi(s)\phi^T(s)}{m^2} ds$$

and we see that the UCO condition is equivalent to the PE condition. □

From the above proof we see another interpretation of the PE property: for a single-output system

$$\begin{aligned} \dot{x} &= 0 \\ y &= c^T(t)x \end{aligned}$$

the PE property of $c(t)$ makes the system observable, i.e., we can recover x from y, even though x is stationary. At any time t, $c^T(t)x$ gives information only about the component of x in the direction of $c(t)$; however, $c(t)$ is moving around in $\mathbb{R}^n$ so that over a finite period of time, we get complete information about x.

$\longrightarrow$ Slow adaptation properties of parameter estimation schemes (adaptive laws) will be useful later for stability analysis of adaptive control systems—this is why we care about things like $\dot{\hat{\theta}} \in L_2$. Also will be relevant for ISS modular design.

Let us take a step back to the static example

$$y(t) = \theta u(t)$$

and recap the gradient method. We have been working with the instantaneous cost, initially defined as

$$J(\hat{\theta}(t), t) := \frac{e^2(t)}{2} = \frac{(\hat{\theta}(t)u(t) - y(t))^2}{2}$$

whose gradient is

$$\nabla J(\hat{\theta}(t), t) = (\hat{\theta}(t)u(t) - y(t))u(t)$$

The graph of J as a function of $\hat{\theta}$ is a parabola which is moving with time (due to the changing input and output). At each t, the minimum is given by

$$\hat{\theta}(t) = \frac{y(t)}{u(t)}$$

Ideally, we'd like to jump to this minimum right away and follow it as it evolves with time. However, we can't really do this. First, it is defined only when $u(t) \neq 0$. Second, this is really sensitive to noise. And in the vector case

$$z(t) = \theta^T(t)\phi(t) \tag{62}$$

we can *never* define $\theta(t)$ by a simple division. So, instead we slide down along the gradient. This takes the data into account over time. It's also clear from this why we need signal boundedness (so that the graph doesn't expand infinitely), and need to work with normalized data if we don't have boundedness.

Now, it would be nice to have a different cost for which we could actually compute the minimum in a well-defined way for each t and follow it as time evolves. Since this doesn't work for the instantaneous cost, we need a cost which, at each time t, takes into account past data up to time t. We can think of this data as a curve (or a set of points) in the u, y plane, and what we're trying to do is really just find a line that best fits this data. Assuming there is no noise, the data would be exactly on a line, and we need to find the slope of this line.

How do we do this?

Lecture 13 (Mon Oct 19)

6.3.3 Least squares

Collect the data into an integral cost, e.g.,

$$J(\hat{\theta}(t), t) := \frac{1}{2}\int_0^t (\hat{\theta}(t)u(s) - y(s))^2 ds$$

(Note that the argument of $\hat{\theta}$ inside the integral is t, not s!) This is still convex in $\hat{\theta}$, and its minimum is where the gradient vanishes:

$$\nabla J(\hat{\theta}) = \int_0^t (\hat{\theta}(t)u(s) - y(s))u(s)ds = \hat{\theta}(t)\int_0^t u^2(s)ds - \int_0^t u(s)y(s)ds$$

which equals 0 when

$$\hat{\theta}(t) = \frac{\int_0^t u(s)y(s)ds}{\int_0^t u^2(s)ds}$$

We see that this is well-defined even if u occasionally equals 0, as long as it is not identically 0 on $[0, t]$. As new data comes, this minimum changes, and it is easy to derive a DE to update it. We'll do this below for the general case.

With a slight modification, we can make the least squares method work for the general case of (62) as well:

$$J(\hat{\theta}) := \frac{1}{2}\int_0^t \underbrace{\frac{(\hat{\theta}^T(t)\phi(s) - z(s))^2}{m^2(s)}}_{(e_n m)^2} ds + \frac{1}{2}(\hat{\theta} - \hat{\theta}_0)^T Q_0 (\hat{\theta} - \hat{\theta}_0)$$

where $Q_0 = Q_0^T > 0$. Here $m^2 = 1 + n_s^2$ as before. The last term penalizes deviations from the initial estimate, and ensures that the least-squares estimate is well defined as we'll see in a minute. The gradient is

$$\begin{aligned}
\nabla J(\hat{\theta}) &= \int_0^t \frac{\hat{\theta}^T(t)\phi(s) - z(s)}{m^2(s)}\phi(s)ds + Q_0(\hat{\theta} - \hat{\theta}_0) \\
&= \left(\int_0^t \frac{\phi(s)\phi^T(s)}{m^2(s)}ds + Q_0\right)\hat{\theta} - \left(\int_0^t \frac{z(s)\phi(s)}{m^2(s)}ds + Q_0\hat{\theta}_0\right)
\end{aligned}$$

so the minimum is at

$$\hat{\theta}(t) = \underbrace{\left(\int_0^t \frac{\phi(s)\phi^T(s)}{m^2(s)}ds + Q_0\right)^{-1}}_{=:\, P(t)} \left(\int_0^t \frac{z(s)\phi(s)}{m^2(s)}ds + Q_0\hat{\theta}_0\right)$$

The matrix $P(t)$ is well defined because $\phi\phi^T \geq 0$ and $Q_0 > 0$ (that's why we needed it).

We want to express this recursively, to avoid computing the inverse and run a DE instead. We have

$$\dot{P} = -P\frac{\phi\phi^T}{m^2}P$$

(derivation: $0 = \frac{d}{dt}(PP^{-1}) = \dot{P}P^{-1} + P\frac{\phi\phi^T}{m^2}$). The differential equation for $\hat{\theta}$ is

$$\dot{\hat{\theta}} = -P\frac{\phi\phi^T}{m^2}\underbrace{P\left(\int_0^t \frac{z(s)\phi(s)}{m^2(s)}ds + Q_0\hat{\theta}_0\right)}_{\hat{\theta}} + P\frac{z\phi}{m^2} = P\phi\frac{z - \hat{\theta}^T\phi}{m^2} = -Pe_n\phi$$

So, the continuous-time recursive least-squares algorithm is:

$$\dot{\hat{\theta}} = -Pe_n\phi$$
$$\dot{P} = -P\frac{\phi\phi^T}{m^2}P$$

This is very similar to the Kalman filter. P is the "covariance matrix" (this is all deterministic, though).

Theorem 7 (Ioannou-Sun, Theorem 4.3.4; Ioannou-Fidan, Theorem 3.7.2)
(i) $\hat{\theta}, P \in L_\infty$ *and* $e_n, e_n n_s, \dot{\hat{\theta}} \in L_2 \cap L_\infty$.
(ii) $\hat{\theta}(t) \to \bar{\theta}$ *where* $\bar{\theta}$ *is some constant vector.*
(iii) If $\frac{\phi}{m}$ *is PE, then* $\bar{\theta} = \theta$.

Item (ii) is a unique property of least squares. On the other hand, note that convergence in (iii) is not claimed to be exponential.

Proof.

(i), (ii) P is positive definite and $\dot{P} \leq 0 \implies P$ is bounded and has a limit. With the usual definition $\tilde{\theta} := \hat{\theta} - \theta$ we have

$$\frac{d}{dt}(P^{-1}\tilde{\theta}) = \frac{\phi\phi^T}{m^2}\tilde{\theta} - P^{-1}P\underbrace{e_n}_{=\frac{\phi^T\tilde{\theta}}{m^2}}\phi = 0 \tag{63}$$

which means that

$$P^{-1}(t)\tilde{\theta}(t) \equiv P^{-1}(0)\tilde{\theta}(0) \quad \implies \quad \tilde{\theta}(t) = P(t)P^{-1}(0)\tilde{\theta}(0)$$

Hence $\tilde{\theta}$ also is bounded and has a limit, and thus so does $\hat{\theta}$. The other claims in (i) are shown using the Lyapunov function

$$V(\tilde{\theta}) := \frac{1}{2}\tilde{\theta}^T P^{-1}\tilde{\theta}$$

in the same way as in the proof of Theorem 5 (do this yourself and check [Ioannou-Sun]).

(iii) We want to show that $P^{-1} \to \infty$ when $\frac{\phi}{m}$ is PE; this would force $\tilde{\theta}$ to go to 0 by (63) and we would be done. Integrating

$$\frac{d}{dt}(P^{-1}) = \frac{\phi\phi^T}{m^2}$$

and using the definition of PE, we have

$$P^{-1}(T_0) - P^{-1}(0) = \int_0^{T_0} \frac{\phi\phi^T}{m^2} ds \geq \alpha_0 T_0 I$$

Similarly,

$$P^{-1}(2T_0) - P^{-1}(T_0) \geq \alpha_0 T_0 I$$

and we see that P^{-1} indeed increases to ∞. □

There are various modifications to the above scheme. For example, one could introduce a "forgetting factor":

$$J(\hat{\theta}) := \frac{1}{2}\int_0^t e^{-\lambda(t-s)} \frac{(\hat{\theta}^T(t)\phi(s) - z(s))^2}{m^2(s)} ds + \frac{1}{2}e^{-\lambda t}(\hat{\theta} - \hat{\theta}_0)^T Q_0 (\hat{\theta} - \hat{\theta}_0)$$

where $\lambda > 0$. Then one can show that convergence of $\hat{\theta}$ to θ in (iii) becomes exponential. The DE for P changes to

$$\dot{P} = \lambda P - P\frac{\phi\phi^T}{m^2}P$$

which actually prevents P from becoming arbitrarily small and slowing down adaptation too much, something that happens with pure least squares where P^{-1} grows without bound. However, items (i) and (ii) will no longer be true.

The above integral cost with forgetting factor (and $Q_0 = 0$) could also be used in the gradient method instead of the instantaneous cost. Although the gradient expression for it is more complicated, the resulting gradient adaptive law has the same properties as the previous one, and in addition one can prove that $\dot{\hat{\theta}} \to 0$. See [Ioannou-Sun, Sect. 4.3.5] for details.

$\longrightarrow$ So, the main difference between the gradient and least squares methods is not so much in the cost used, but in the method itself (moving along the negative gradient of a time-varying function, versus following its exact minimum).

6.3.4 Projection

In the above, we were assuming that $\theta \in \mathbb{R}^m$ is arbitrary, i.e., completely unknown. In practice, we usually have some set S which we know contains θ, and we can take this knowledge into account when designing an adaptive law. The basic idea is to use projection to ensure that $\hat{\theta}(t) \in S \; \forall t$. This is a good idea for two reasons:

- Improves convergence and transient behavior of the estimation scheme.
- Helps avoid estimated plant models which are not suitable for control purposes; e.g, for some values of $\hat{\theta}$ the model may not be stabilizable (*loss of stabilizability* problem) and certainty equivalence-based design breaks down. (More on this later.)

For example, suppose that we have some cost $J(\hat{\theta})$ and want to implement the gradient law

$$\dot{\hat{\theta}} = -\nabla J(\hat{\theta})$$

but want to use the fact that $\theta \in S$ where S is given by

$$S = \{\theta : g(\theta) = 0\}$$

for some (nice) scalar-valued function g. This captures cases where θ belongs to some known subspace. Given a $\hat{\theta} \in S$, we can split $-\nabla J(\hat{\theta})$ into a sum of two terms, one normal to S and the other tangent to S:

$$-\nabla J(\hat{\theta}) = \alpha \underbrace{\nabla g(\hat{\theta})}_{\text{normal}} + \underbrace{\Pr(\hat{\theta})}_{\text{tangent}} \tag{64}$$

Then we will discard the normal direction and just keep the tangent direction. This will ensure that $\hat{\theta}$ travels along S.

But how do we calculate $\Pr(\hat{\theta})$? For this we need to know α. Multiply both sides of (64) by the normal, i.e., $(\nabla g(\hat{\theta}))^T$:

$$-\nabla g^T \nabla J = \alpha \nabla g^T \nabla g \quad \Longrightarrow \quad \alpha = -\frac{\nabla g^T \nabla J}{\nabla g^T \nabla g}$$

and so

$$\Pr = -\nabla J + \frac{\nabla g^T \nabla J}{\nabla g^T \nabla g} \nabla g = -\left(I - \frac{\nabla g \nabla g^T}{\nabla g^T \nabla g}\right) \nabla J$$

This gives us the projected gradient law.

More generally, if we started with the scaled gradient law

$$\dot{\hat{\theta}} = -\Gamma \nabla J(\hat{\theta})$$

(which is a coordinate change away from the previous one) then its projected version can be shown to be

$$\dot{\hat{\theta}} = -\left(I - \Gamma \frac{\nabla g \nabla g^T}{\nabla g^T \Gamma \nabla g}\right) \Gamma \nabla J(\hat{\theta})$$

It is easy to extend this idea to the case when the set S is given by an inequality constraint:

$$S = \{\theta : g(\theta) \leq 0\}$$

For example, S could be a ball. Start with $\hat{\theta}_0 \in S$. If $\hat{\theta}(t)$ is in the interior of S, or it's on the boundary but $-\Gamma\nabla J(\hat{\theta})$ points inside S, then follow the usual gradient method. If $\hat{\theta}(t)$ is on the boundary and $-\Gamma\nabla J(\hat{\theta})$ points outside, then apply the projection.

$\longrightarrow$ If S is a convex set, then the projected gradient adaptive law retains all properties of the unprojected gradient adaptive law established in Theorem 5 (and in addition maintains $\hat{\theta}$ in S by construction).

This is true because when we subtract a term proportional to $\nabla g(\hat{\theta})$ from $\dot{\hat{\theta}}$, this can only make the Lyapunov function (59) decrease more. When the sublevel sets of g are convex, the contribution of $-\nabla g(\hat{\theta})$ to $\dot{V}$ is negative basically for the same reason that the contribution of $-\nabla J(\hat{\theta})$ was negative for the convex cost function J. See [Ioannou-Sun, Sect. 4.4] for details.

The projection modification can also be applied to the least squares method.

6.4 Sufficiently rich signals and parameter identification

We've been studying the parametric model

$$z = \theta^T \phi \tag{65}$$

and have developed parameter estimation schemes which guarantee parameter convergence ($\hat{\theta} \to \theta$, exponentially except for pure least squares, need forgetting factor) if ϕ is PE. The regressor vector ϕ contains (filtered version of) u, y, and their derivatives, see (55). We can write this as

$$\phi = H(s)u$$

But the only signal we have direct control over is u.

Basic idea:

$$u \text{ sufficiently rich} \quad \Longrightarrow \quad \phi \text{ is PE} \quad \Longrightarrow \quad \hat{\theta}(t) \to \theta \tag{66}$$

We want to understand what "sufficiently rich" means in various cases.

Example 7 (we already had this earlier) Scalar plant:

$$\dot{y} = -ay + bu$$

Assume $a > 0$ and $u \in L_\infty$. In the parametric model for this example, we have

$$z = \dot{y}, \qquad \theta = \begin{pmatrix} b \\ a \end{pmatrix}, \qquad \phi = \begin{pmatrix} u \\ -y \end{pmatrix} = \begin{pmatrix} 1 \\ -\frac{b}{s+a} \end{pmatrix} u$$

(we ignore the filtering).

What would be a good input to inject?

Let's try a constant input $u = c$:

$$\dot{y} = -ay + bc = -a\Big(y - \frac{bc}{a}\Big)$$

hence

$$y(t) = \frac{bc}{a} + e^{-at}\Big(y_0 - \frac{bc}{a}\Big)$$

Can we recover both parameters a and b by observing y?

No. The easiest way to see this is to consider the case when $y_0 = \dfrac{bc}{a}$, then $y(t) = \dfrac{bc}{a}$ and this only carries information about the ratio $\dfrac{b}{a}$.

For other initial conditions, we do have information on a itself but it disappears exponentially fast and in steady state y only carries information about $\dfrac{b}{a}$.

Accordingly, we can show that

$$\phi = \begin{pmatrix} c \\ -\frac{bc}{a} - e^{-at}\Big(y_0 - \frac{bc}{a}\Big) \end{pmatrix}$$

is not PE. Think about it: when we form the matrix $\phi(t)\phi^T(t)$, it is the sum of two terms. One is constant and singular. The other decays exponentially to 0. So the definition of PE cannot be satisfied.

What would be a better input? How does one in general identify frequency response of a system with transfer function $g(s)$? (In our case, $g(s) = \frac{b}{s+a}$).

Let's try

$$u(t) = \sin \omega t$$

where ω is some constant frequency. In steady state (i.e., modulo an exponentially decaying transient), the corresponding output is

$$y(t) = \operatorname{Re} g(j\omega) \sin \omega t + \operatorname{Im} g(j\omega) \cos \omega t = A \sin(\omega t + \alpha)$$

where

$$\begin{aligned} A &:= |g(j\omega)| = \left|\frac{b}{j\omega + a}\right| = \frac{|b|}{\sqrt{\omega^2 + a^2}}, \\ \alpha &:= \tan^{-1} \frac{\operatorname{Im} g(j\omega)}{\operatorname{Re} g(j\omega)} = -\tan^{-1} \frac{\omega}{a} \end{aligned} \tag{67}$$

where the second formula is valid if $b > 0$, otherwise add $180°$ to α.

By observing y we can get A and α from which, using algebraic operations, we can recover a and b. (Note that now the information doesn't disappear with time).

Accordingly,

$$\phi = \begin{pmatrix} \sin\omega t \\ -A\sin(\omega t + \alpha) \end{pmatrix}$$

can be shown to be PE; see homework. □

Example 8 Second-order plant:

$$y(s) = \frac{b_1 s + b_0}{s^2 + a_1 s + a_0} u(s)$$

where $a_0, a_1 > 0$ to ensure stability. There are 4 unknown parameters in the transfer function now, so a single sinusoid is not enough. But

$$u(t) = \sin\omega_1 t + \sin\omega_2 t, \qquad \omega_1 \neq \omega_2$$

gives

$$y(t) = A_1 \sin(\omega_1 t + \alpha_1) + A_2 \sin(\omega_2 t + \alpha_2)$$

We now have 4 unknowns and 4 equations $\implies$ OK. □

Lecture 14 (Wed Oct 21)

Back to the general parametric model (65) where θ and ϕ both have dimension m.

Generically, u is sufficiently rich (in the sense of guaranteeing that ϕ is PE) if it contains at least $\frac{m}{2}$ distinct frequencies. This matches the above examples.

The above definition assumes that u is sinusoidal. We can make this more general, by saying that the power spectrum contains at least m distinct points (power spectrum is symmetric).

"Generically" means modulo some degenerate situations, e.g., when frequencies ω_i coincide with zeros of $H(s)$.

The proof (which we will not give) relies on frequency domain arguments. The basic idea is that the integral appearing in the definition of PE is the autocovariance of ϕ, and it is related to the power spectrum (or spectral measure) of ϕ via Fourier transform. See [Ioannou-Sun, Sect. 5.2.1] for details.

Note: if the plant is partially known, then we need fewer distinct frequencies. Or, with fewer frequencies we can partially identify the unknown plant.

Earlier, we described parameter estimation schemes of two kinds: using state-space models, and using parametric models. Combining them with the discussion we just had, we can formally state results on parameter convergence. Note that there are two steps in (66). The above discussion was primarily about the first step (we did talk about recovering the system parameters but this was only for simple examples, in general we need adaptive laws to do that). The second step was addressed—for parametric models only—in Theorems 5 and 7.

For full-state measurements, we had

$$\dot{x} = Ax + Bu$$

where A is Hurwitz, $x \in \mathbb{R}^n$, $u \in L_\infty$. For simplicity assume that u is scalar (the multi-input case is in [Ioannou-Sun]). The adaptive law design was given on pages 55–56, but we didn't study parameter convergence. We can now state the following property of that scheme.

Theorem 8 (Ioannou-Sun, Theorem 5.2.2) *If (A, B) is a controllable pair and u contains at least $\frac{n+1}{2}$ distinct frequencies, then the adaptive law discussed in class gives $\hat{A} \to A$, $\hat{B} \to B$ (exponentially fast).*

We don't give a proof, but note that the number $n+1$ comes from the fact that state dimension is n plus we have 1 input. If there are uncontrollable modes, then they are not affected by u and decay to 0, so the corresponding parameters cannot be identified.

For partial state measurements, we had

$$\begin{aligned} \dot{x} &= Ax + bu \\ y &= c^T x \end{aligned}$$

where A is Hurwitz, $x \in \mathbb{R}^n$, and $u, y \in \mathbb{R}$ are bounded. We already said that, even though the total number of parameters in the state model is $n^2 + 2n$, the dimension of the parametric model (the number of unknown coefficients in the transfer function) is at most $m = 2n$. We can only hope to identify these $2n$ parameters. We have described adaptive laws based on the gradient method and least-square method. Since we're in the stable case here, we don't need normalization.

Theorem 9 (Ioannou-Sun, Theorem 5.2.4) *If the transfer function $g(s)$ has no pole-zero cancellations and u contains at least n distinct frequencies, then the adaptive laws discussed in class give $\hat{\theta} \to 0$ (exponentially fast, except for the case of pure least squares).*

This result matches the examples we had earlier. As for the previous result, we need fewer frequencies there because we are observing more information (the whole state instead of output).

Some remarks on parameter estimation/identification:

- As we emphasized several times, parameter identification is not necessary for satisfactory control (already saw in the first example).

- When considering control objectives such as stabilization, PE is actually in direct conflict with control, since we don't want to use inputs such as sinusoids.

- Even without PE assumptions, we know that parameter estimation schemes have useful properties. Look, for example, at statement (i) of Theorem 5, and similar statements in other results we had. These properties are directly useful in adaptive control, as we will see. PE conditions, on the other hand, are not of primary importance for this course.

- There are, however, other control objectives such as reference tracking or model matching. For these, satisfying PE conditions may sometimes be feasible.

- Adaptive control schemes based on estimating unknown plant parameters, and then using them for control, are called *indirect*. In contrast, in *direct* adaptive control one works directly with controller parameters: first, reparameterize the problem in terms of unknown desired controller parameters, and then design an adaptive law for estimating them. (See the block diagrams in [Ioannou-Sun, pp. 10–11].)

The next example, on *model reference adaptive control*, is taken from [Ioannou-Sun, pp. 320 and 397] and [Khalil, pp. 16 and 327]. It illustrates the last two points above. It also reinforces the concept of PE and other constructions we saw above. This example will be used again later when we study singular perturbations. See [Ioannous-Sun, Chap. 6] for more information on model-reference adaptive control.

$\longrightarrow$ This will also be the first time we combine parameter estimation with control (not counting our early example which wasn't really using an estimation scheme but an ad-hoc tuning law).

6.5 Case study: model reference adaptive control

Model reference adaptive control (MRAC): want the closed-loop system to reproduce the input/output behavior of a given reference model, driven by the reference signal r.

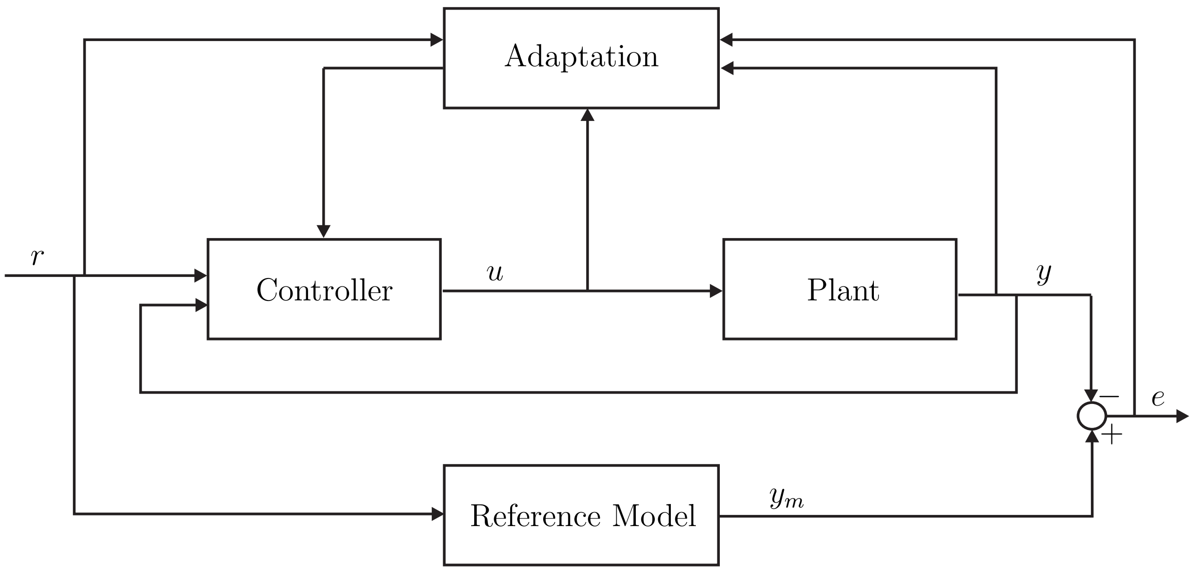

Consider the scalar plant

$$\dot{y} = ay + bu$$

Here a is completely unknown, i.e., we don't assume that we know its sign—hence we no longer use the minus sign we had earlier. The control gain b is also unknown but we assume that we know its sign; for concreteness, suppose $b > 0$.

Reference model:

$$\dot{y}_m = -a_m y_m + b_m r \tag{68}$$

We assume that $a_m > 0$ (reference model is stable) and $r \in L_\infty$.

For control, we use a combination of feedforward and feedback of the form

$$u = -ky + lr$$

which gives

$$\dot{y} = (a - bk)y + blr$$

Hence to match the model, we need

$$a - bk = -a_m, \quad bl = b_m \quad \Leftrightarrow \quad k = \frac{a + a_m}{b}, \quad l = \frac{b_m}{b}$$

These controller gains k, l are not implementable, since we don't know a and b. Instead, we will use

$$u = -\hat{k}y + \hat{l}r \tag{69}$$

where $\hat{k}, \hat{l}$ are estimates of k, l. There are two approaches to generating such an adaptive control law: direct and indirect.

6.5.1 Direct MRAC

Let us rewrite the plant as

$$\dot{y} = -a_m y + b_m r + b(u + ky - lr) \tag{70}$$

This is a reparameterization which directly displays control parameters; note that it is bilinear in the parameters, not linear. To check (70), just plug in the formulas for k and l.

Plugging the control law (69) into (70), we have

$$\dot{y} = -a_m y + b_m r + b(-\tilde{k}y + \tilde{l}r) \tag{71}$$

where we defined the parameter errors

$$\tilde{k} := \hat{k} - k, \qquad \tilde{l} := \hat{l} - l$$

Define also the output tracking error

$$e := y_m - y$$

Using (68) and (71), we compute:

$$\dot{e} = -a_m e + b\tilde{k}y - b\tilde{l}r$$

Let's try the candidate Lyapunov function

$$V(e, \tilde{k}, \tilde{l}) := \frac{1}{2}\left(\frac{e^2}{b} + \frac{1}{\gamma}(\tilde{k}^2 + \tilde{l}^2)\right)$$

where $\gamma > 0$ is arbitrary.

$\longrightarrow$ The reason for division by b will be clear soon; recall that we assumed $b > 0$.

$\longrightarrow$ Division by γ is just to have an extra degree of freedom; γ will play the same role in the adaptive law as the scaling factor in the gradient method. (We didn't use it in Lyapunov-based design earlier, e.g., in (52), but we could have used it there.) We could be even more general and divide $\tilde{k}^2$ by γ_1 and $\tilde{l}^2$ by $\gamma_2 \neq \gamma_1$.

The derivative of V is

$$\dot{V} = -\frac{a_m}{b}e^2 + \tilde{k}ey - \tilde{l}er + \frac{1}{\gamma}\tilde{k}\dot{\tilde{k}} + \frac{1}{\gamma}\tilde{l}\dot{\tilde{l}}$$

where $\dot{\tilde{k}} = \dot{\hat{k}}$ and $\dot{\tilde{l}} = \dot{\hat{l}}$ are yet to be specified. But now their choice is clear:

$$\dot{\hat{k}} := -\gamma e y, \qquad \dot{\hat{l}} := \gamma e r$$

which gives

$$\dot{V} = -\frac{a_m}{b} e^2 \tag{72}$$

Now the analysis follows the familiar steps:

$e, \tilde{k}, \tilde{l}$ are bounded.

$e \in L_2$.

We assumed that r is bounded and the reference model is stable. Hence y_m is bounded, which since e is bounded implies that $y = y_m - e$ is bounded. Hence, $\dot{e}$ is bounded.

By Barbalat, $e \to 0$. Also, $\dot{\hat{k}}, \dot{\hat{l}}$ are in L_2 and converge to 0.

We showed that, for any bounded input, the plant output y asymptotically tracks the reference model output y_m, which was the original objective.

Do we know that $\hat{k}, \hat{l}$ converge to k, l?

Need a PE condition. More precisely, we need the reference signal r to be sufficiently rich. Since the plant is scalar, we guess that one sinusoid will do:

$$r = \sin \omega t$$

which *in steady state* gives the output[8]

$$y = A \sin(\omega t + \alpha)$$

where A and α are given in (67). Collecting the above equations, we can write the resulting system in $(e, \tilde{k}, \tilde{l})$ coordinates as an LTV system:

$$\begin{pmatrix} \dot{e} \\ \dot{\tilde{k}} \\ \dot{\tilde{l}} \end{pmatrix} = \begin{pmatrix} -a_m & bA\sin(\omega t + \alpha) & -b \sin \omega t \\ -\gamma A \sin(\omega t + \alpha) & 0 & 0 \\ \gamma \sin \omega t & 0 & 0 \end{pmatrix} \begin{pmatrix} e \\ \tilde{k} \\ \tilde{l} \end{pmatrix} =: A(t) \begin{pmatrix} e \\ \tilde{k} \\ \tilde{l} \end{pmatrix}$$

Looking at (72), we see that we'll have exponential stability of this system—and thus exponential convergence of the controller parameter estimates $\hat{k}$ and $\hat{l}$ to their true values—if we can show that the system is UCO with respect to the output

$$y := e = \begin{pmatrix} 1 & 0 & 0 \end{pmatrix} \begin{pmatrix} e \\ \tilde{k} \\ \tilde{l} \end{pmatrix} =: C \begin{pmatrix} e \\ \tilde{k} \\ \tilde{l} \end{pmatrix}$$

[8] Note that, first, the reference model is stable so the effects of its initial conditions vanish with time, and second, the steady-state responses of the plant and of the reference model are the same because $y_m - y = e \to 0$. So, what we're doing is ignore some terms that go to 0 with time, and it can be shown that the presence of these terms doesn't affect stability which we're about to establish without these terms (cf. [Khalil, Example 9.6].

We also know the trick of passing from $(A(t), C)$ to $(A(t) + L(t)C, C)$, see Lemma 6. It is easy to see that by proper choice of $L(t)$, we can get

$$A(t) + L(t)C = \begin{pmatrix} -a_m & bA\sin(\omega t + \alpha) & -b\sin\omega t \\ 0 & 0 & 0 \\ 0 & 0 & 0 \end{pmatrix}$$

(we could also kill $-a_m$ if we wanted, but it helps to keep it). Still, showing UCO (defined in Section 2.2) looks complicated because we need to compute the transition matrix of $A(t) + L(t)C$, and it's not clear how to do that.

$\longrightarrow$ By the way, it is straightforward to check that in case r were constant instead of sinusoidal, and hence y were also constant in steady state, the above (LTI) pair would not be observable.

Consider the system $\dot{x} = (A + LC)x$, $y = Cx$ (note: this y is not to be confused with the original plant output y, we are considering an abstract system here). To investigate its UCO, note that the observability Gramian $M(t_0, t_0 + T)$ by definition satisfies

$$\int_{t_0}^{t_0+T} |y(t)|^2 dt = x(t_0)^T M(t_0, t_0 + T)x(t_0)$$

hence all we need to show is that

$$\int_{t_0}^{t_0+T} |y(t)|^2 dt \geq \alpha_1 |x(t_0)|^2 \tag{73}$$

for some $\alpha_1 > 0$. But we have $y = x_1$ and

$$\dot{y} = -a_m y + bA\sin(\omega t + \alpha)x_2 - b\sin\omega t x_3$$

and also

$$\dot{x}_2 = \dot{x}_3 = 0 \quad \Longrightarrow \quad x_2, x_3 = \text{const}$$

(note that x_2, x_3 no longer correspond to $\tilde{k}, \tilde{l}$ because we're working with an auxiliary system obtained by output injection).

We can now easily solve for y:

$$y(t) = e^{-a_m t}y(t_0) + \begin{pmatrix} \int_{t_0}^t e^{-a_m(t-s)}bA\sin(\omega s + \alpha)ds & -\int_{t_0}^t e^{-a_m(t-s)}b\sin\omega s ds \end{pmatrix} \begin{pmatrix} x_2 \\ x_3 \end{pmatrix}$$

We know that the vector

$$\begin{pmatrix} A\sin(\omega t + \alpha) \\ \sin\omega t \end{pmatrix}$$

is PE. The integral terms are outputs of the stable linear system $\dot{y} = -a_m y + u$ driven (componentwise) by this PE input vector, and thus can be shown to form a PE vector as well. With some calculations, the desired property (73) follows. For details, see [Ioannou-Sun, Lemmas 4.8.3 and 4.8.4].

So, the conclusion is that output tracking is achieved for any bounded reference, and in addition for sinusoidal references the controller parameters are identified.

Lecture 15 (Mon Oct 26)

6.5.2 Indirect MRAC

We still want to use a control law of the form

$$u = -\hat{k}y + \hat{l}r \tag{74}$$

but instead of defining $\hat{k}, \hat{l}$ directly as estimates of k and l, we'll now define them indirectly via estimates $\hat{a}, \hat{b}$ of the plant parameters a, b. Earlier we derived the ideal control parameters to be

$$k = \frac{a + a_m}{b}, \qquad l = \frac{b_m}{b}$$

so we set

$$\hat{k} := \frac{\hat{a} + a_m}{\hat{b}}, \qquad \hat{l} = \frac{b_m}{\hat{b}} \tag{75}$$

To generate $\hat{a}$ and $\hat{b}$, we can follow a familiar scheme discussed earlier—cf. Example 6 in Section 6.2.

Comparing with Example 6, we know that we would like to run the estimator

$$\dot{\hat{y}} = -a_m(\hat{y} - y) + \hat{a}y + \hat{b}u$$

Note that we choose the damping rate of the estimator to be the same as in the reference model. As a consequence, it turns out that we don't actually need to run the estimator, and can just use the state y_m of the reference model instead of $\hat{y}$. Indeed, let's rewrite the reference model like this:

$$\begin{aligned} \dot{y}_m &= -a_m y_m + b_m r = -a_m(y_m - y) - a_m y + b_m r = -a_m(y_m - y) + \hat{a}y - (\hat{a} + a_m)y + b_m r \\ &= -a_m(y_m - y) + \hat{a}y + \hat{b}\underbrace{\left(-\frac{\hat{a} + a_m}{\hat{b}}y + \frac{b_m}{\hat{b}}r\right)}_{=u} = -a_m(y_m - y) + \hat{a}y + \hat{b}u \end{aligned} \tag{76}$$

(this is valid for the closed-loop system).

The tracking error is

$$e := y_m - y$$

and its derivative is (using the final equation for $\dot{y}_m$ above)

$$\dot{e} = \dot{y}_m - \dot{y} = -a_m e + \tilde{a}y + \tilde{b}u \tag{77}$$

where, as usual,

$$\tilde{a} := \hat{a} - a, \qquad \tilde{b} := \hat{b} - b$$

The next step is Lyapunov analysis.

Candidate Lyapunov function:

$$V(e, \tilde{a}, \tilde{b}) := \frac{1}{2}\left(e^2 + \frac{1}{\gamma}(\tilde{a}^2 + \tilde{b}^2)\right)$$

where $\gamma > 0$ is arbitrary. This is the same V as we've been using earlier, with the parameter γ for extra flexibility. Note that we don't need the division by b as in the direct scheme, since we're no longer using the bilinear parameterization we used there and so b doesn't multiply the parameter errors in the $\dot{e}$ equation. (In fact, we're not seeing the need for assuming $b > 0$ yet, but we'll see it later.)

$$\dot{V} = -a_m e^2 + \tilde{a}ey + \tilde{b}eu + \frac{1}{\gamma}\tilde{a}\dot{\tilde{a}} + \frac{1}{\gamma}\tilde{b}\dot{\tilde{b}}$$

which makes the choice of adaptive law, i.e., the choice of $\dot{\tilde{a}} = \dot{\hat{a}}$ and $\dot{\tilde{b}} = \dot{\hat{b}}$, obvious:

$$\dot{\hat{a}} := -\gamma ey, \qquad \dot{\hat{b}} = -\gamma eu$$

This is very similar to what we had in the direct scheme, and it gives

$$\dot{V} = -a_m e^2$$

Following the standard steps, we get $e, \hat{a}, \hat{b} \in L_\infty$, $e \in L_2$.

Need to look at $\dot{e}$, given by (77). If we can show that it's bounded, then we'll have $e \to 0$ by Barbalat, meaning that asymptotic tracking is achieved.

As before, since we assumed that r is bounded the reference model is stable, y_m is bounded, hence $y = y_m - e$ is bounded.

But we also need to know that u is bounded. It is given by (74), (75).

$\longrightarrow$ Here comes the major difference with the direct case: unlike in the direct case, the Lyapunov analysis doesn't tell us anything about boundedness of $\hat{k}$ and $\hat{l}$.

Is boundedness of u guaranteed in the present scheme?

In general, no, because there's nothing to prevent $\hat{b}$ from becoming arbitrarily close to 0, or even hit 0. Then u will become unbounded, or will not even be defined.

To understand why $\hat{b} = 0$ is a problem, think about the plant model

$$\dot{y} = \hat{a}y + \hat{b}u$$

on which our certainty equivalence controller is based. When $\hat{b} = 0$, the model is not stabilizable, and the procedure breaks down (u doesn't exist). This is a well-known issue in indirect adaptive control, known as the *loss of stabilizability problem.*

There are several approaches for dealing with loss of stabilizability; see [Ioannou-Sun, Sect. 7.6.2]. The simplest is based on *projection* and works if we have a bit more information on the plant, as follows. We assumed that $b > 0$, i.e., we know that the actual plant is stabilizable. Suppose that we know a constant $b_0 > 0$ such that

$$b \geq b_0 \text{ (known)}$$

Then, we can modify the adaptive law for $\hat{b}$ by projecting it onto the interval $[b_0, \infty)$:

$$\dot{\hat{b}} = \begin{cases} -\gamma e u, & \text{if } \hat{b} > b_0 \text{ or if } \hat{b} = b_0 \text{ and } eu \leq 0 \\ 0, & \text{otherwise} \end{cases}$$

(of course, we initialize $\hat{b}$ at a value higher than b_0).

The analysis should be slightly modified, because when $\hat{b}$ stops, from the $\dot{V}$ equation we have

$$\dot{V} = -a_m e^2 + \tilde{b} e u$$

However, when we stop, we know by construction that $\hat{b} = b_0$ and $eu > 0$. Since $b \geq b_0$, at this time we have $\tilde{b} = \hat{b} - b \leq 0$. Hence, the extra term in $\dot{V}$ is nonpositive, and we're still OK. And now $\hat{b}$ is bounded away from 0, hence $\hat{k}, \hat{l}$ are bounded, u is bounded, $\dot{e}$ is bounded, and we can apply Barbalat to conclude that $e \to 0$ and tracking is achieved.

Also, $\dot{\hat{a}}, \dot{\hat{b}}$ are in L_2 and converge to 0.

$\longrightarrow$ Note that $\hat{b}$ may or may not approach 0 in practice, we just can't rule it out theoretically in the absence of projection. In this example the loss of stabilizability issue is not really severe because we know the sign of b. In problems where the sign of b is unknown (as we had earlier) and thus $\hat{b}$ might be initialized with the wrong sign, zero crossings for gradient adaptive laws are in fact almost unavoidable. But in higher dimensions things are different.

$\longrightarrow$ What we proved about exponential parameter convergence under sinusoidal reference signals for direct MRAC is still true for indirect MRAC.

From this example, it may appear that direct MRAC is better because it doesn't suffer from the loss of stabilizability issue. However, in general this is not the case. The reason is that direct MRAC requires us to come up with a controller reparameterization which can then be used to design an adaptive law. In the above simple example this worked, although the direct controller reparameterization was more complicated than the original plant parameterization (it was bilinear). In general, direct schemes tend to be more complicated and apply to narrower classes of systems than do indirect schemes. (This is true not only for MRAC.) See [Ioannou-Sun, Sect. 1.2.3] for a detailed discussion. In this course, we're primarily dealing with indirect adaptive control.

Note that the estimation scheme we used for indirect MRAC is quite different from the ones we developed in Section 6.3, in several key aspects:

- It does not use normalization.
- It relies on an estimator/reference model (if we don't ask for parameter identification, then we can set $r = 0$ if we want) instead of plant parametric model.

- It guarantees that $e \to 0$, something we didn't have with the normalized adaptive laws based on parametric models.
- Since the design relies on the matching between the estimator and the reference model, given by (76), it is difficult to extend it to higher-order plants (actually, the issue is not so much the plant order but its *relative degree*), while the earlier schemes relying on parametric models work for general plants.

See [Ioannou-Sun, Chapter 6] for more information.

7 Input-to-state stability

7.1 Weakness of certainty equivalence

Reference: [KKK book, Sect. 5.1]

When using the certainty equivalence principle for control design, we substitute $\hat{\theta}$ for θ in the controller. The resulting closed-loop system thus differs from the "ideal" one by terms involving the parameter estimation error $\tilde{\theta}$. We then hope that these terms do not have too much negative effect on system stability. However, we haven't really formally addressed this issue yet.

Ideally, we'd like to have *modular design*, i.e., formulate some design objectives for the controller and the tuning law (tuner) separately so that, when combined using certainty equivalence, they give stability. For example:

$$\left.\begin{array}{l}\text{Controller is stabilizing when } \tilde{\theta} = 0 \\ \text{Tuner guarantees bounded } \tilde{\theta}\end{array}\right\} \implies x \text{ is bounded} \tag{78}$$

Another conjecture might be:

$$\left.\begin{array}{l}\text{Controller is stabilizing when } \tilde{\theta} = 0 \\ \text{Tuner guarantees } \tilde{\theta} \to 0\end{array}\right\} \implies x \to 0 \tag{79}$$

In our old Example 1, we had the plant

$$\dot{x} = \theta x + u$$

and the controller

$$u = -x - \hat{\theta} x$$

which give the closed-loop dynamics

$$\dot{x} = -x - \tilde{\theta} x$$

Let's view the error $\tilde{\theta}$ as a (disturbance) input to this system. We've shown earlier that it is bounded (but doesn't necessarily converge to 0).

Is it true that a bounded $\tilde{\theta}$ always leads to a bounded x for the above system?

No. For example, if $\tilde{\theta} \equiv -2$ then we have $\dot{x} = x$ and $x \to \infty$.

The tuning law that we used for Example 1, $\dot{\hat{\theta}} = x^2$, provides more than just boundedness of $\tilde{\theta}$. In fact, the Lyapunov analysis we had before implies that $x \to 0$. Some other tuning law which guarantees boundedness of $\tilde{\theta}$ but doesn't ensure its correct sign may not work, so we have to be careful.

So the first property above, (78), is *not true in general*, as the above example shows.

What about the second one, (79)?

In the above example, it is clear that if $\tilde{\theta} \to 0$ then $x \to 0$ (regardless of the sign of $\tilde{\theta}$). But the above plant is linear. The next example shows that for *nonlinear plants*, much more drastic behavior is possible: x may not converge to 0 and *may not even be bounded* despite the fact that $\tilde{\theta} \to 0$!

Example 9 Consider the scalar plant

$$\dot{x} = -\theta x^2 + u$$

An obvious choice for the controller is

$$u = -x + \hat{\theta} x^2$$

and the closed-loop system is

$$\dot{x} = -x + \tilde{\theta} x^2$$

Suppose that we were able to design a parameter identifier which gives

$$\tilde{\theta}(t) = e^{-t}\tilde{\theta}_0$$

i.e., $\tilde{\theta} \to 0$ exponentially fast (we take the decay rate to be equal to 1 just for simplicity, this is not important). This seems to be as good as it gets.

$\longrightarrow$ We haven't discussed parameter identification for nonlinear plants, so we just take such an identifier as given and don't worry about actually designing it. One example of such identifier design is discussed in [KKK book, p. 187], see also the references cited there.

The closed-loop system is now

$$\dot{x} = -x + e^{-t}\tilde{\theta}_0 x^2 \tag{80}$$

What can we then say about the closed-loop system (80)? Is x bounded? Does it converge to 0?

Claim: For some initial conditions, x escapes to ∞ in finite time!

One system that's known to do that is

$$\dot{x} = x^2$$

whose solution is

$$x(t) = \frac{x_0}{1 - x_0 t}$$

and for $x_0 > 0$ this is only defined on the finite time interval $[0, 1/x_0)$ and approaches ∞ as $t \to 1/x_0$. This behavior is due to the rapid nonlinear growth at infinity of the function $f(x) = x^2$.

It is clear that the system

$$\dot{x} = -x + 2x^2 = -x + x^2 + x^2 \tag{81}$$

also has finite escape time, for x_0 large enough, because for large positive x the term $-x$ is dominated by one of the x^2 terms. (This is despite the fact that locally around 0 this system is asymptotically stable.)

But now we can also see that our closed-loop system (80) has finite escape time. The argument is like this: let T be the time that it takes for the solution of (81) to escape to ∞ for a given $x_0 > 0$. Choose $\tilde{\theta}_0$ large enough so that

$$e^{-t}\tilde{\theta}_0 \geq 2 \qquad \forall\, t \in [0, T]$$

Then the corresponding solution of (80) is no smaller than that of (81) and hence it escapes to ∞ in time $\leq T$.

In fact, one can check that the solution of (80) is given by the formula

$$x(t) = x_0 \frac{2}{x_0\tilde{\theta}_0 e^{-t} + (2 - x_0\tilde{\theta}_0)e^t}$$

and we see that for

$$x_0\tilde{\theta}_0 > 2$$

we have

$$x(t) \to \infty \quad \text{as} \quad t \to \frac{1}{2} \log \frac{x_0\tilde{\theta}_0}{x_0\tilde{\theta}_0 - 2}$$

□

The above example highlights the challenge of nonlinear adaptive control and weakness of certainty equivalence approach: bounded/converging estimation error does not guarantee boundedness/convergence of the closed-loop system.

For a modular design to work, we need to demand more from the controller. Namely, we need the controller to possess some *robustness* with respect to disturbance inputs which in our case correspond to the parameter estimation error $\tilde{\theta}$. Such robustness properties are captured in nonlinear system theory by the general concept of *input-to-state stability (ISS)* which is discussed next.

Lecture 16 (Wed Oct 28)

7.2 Input-to-state stability and stabilization

References: [Khalil, Sect. 4.9]; [KKK book, Appendix C]; papers by Eduardo Sontag, in particular his recent survey paper "Input to state stability: basic concepts and results" downloadable from his website, as well as the paper "Universal construction of feedback laws achieving ISS and integral-ISS disturbance attenuation" by Sontag, Wang and myself, Systems and Control Letters, vol. 4, pp. 111-127, 2002.

First, we need to define a few useful function classes.

A function $\gamma : [0, \infty) \to [0, \infty)$ is said to be of *class* $\mathcal{K}$ if it is continuous, strictly increasing, and $\gamma(0) = 0$.

If γ is also unbounded, then it is said to be of *class* $\mathcal{K}_\infty$. Example: $\gamma(r) = cr$ for some $c > 0$.

A function $\beta : [0, \infty) \times [0, \infty) \to [0, \infty)$ is said to be of *class* $\mathcal{KL}$ if $\beta(\cdot, t)$ is of class $\mathcal{K}$ for each fixed $t \geq 0$ and $\beta(r, t)$ is decreasing to zero as $t \to \infty$ for each fixed $r \geq 0$. Example: $\beta(r, t) = ce^{-\lambda t}r$ for some $c, \lambda > 0$.

We will write $\beta \in \mathcal{KL}$, $\gamma \in \mathcal{K}_\infty$ to indicate that β is a class $\mathcal{KL}$ function and γ is a class $\mathcal{K}_\infty$ function, respectively.

Now, consider a general nonlinear system

$$\dot{x} = f(x, d) \tag{82}$$

where $x \in \mathbb{R}^n$ is the state and $d \in \mathbb{R}^\ell$ is the (disturbance) input. To ensure existence and uniqueness of solutions, we assume that f is sufficiently nice (e.g., locally Lipschitz) and d is also sufficiently nice (e.g., piecewise continuous).

According to [Sontag, 1989] the system (82) is called *input-to-state stable (ISS)* with respect to d if for some functions $\beta \in \mathcal{KL}$ and $\gamma \in \mathcal{K}_\infty$, for every initial state x_0, and every input d the corresponding solution of (82) satisfies the inequality

$$|x(t)| \leq \beta(|x_0|, t) + \gamma\big(\sup_{0 \leq s \leq t} |d(s)|\big) \qquad \forall\, t \geq 0$$

The above formula assumes that the initial time is $t_0 = 0$. But since the system (82) is time-invariant, it wouldn't make any difference if we kept the initial time t_0 general and wrote

$$|x(t)| \leq \beta(|x_0|, t - t_0) + \gamma\big(\sup_{t_0 \leq s \leq t} |d(s)|\big) \qquad \forall\, t \geq t_0 \geq 0. \tag{83}$$

$\longrightarrow$ Also note that by causality, it makes no difference if we take sup over $t_0 \leq s < \infty$.

Let's try to decipher what the ISS definition says. When there is no input, i.e., when $d \equiv 0$, it reduces to

$$|x(t)| \leq \beta(|x_0|, t)$$

This means that the state is upper-bounded by $\beta(|x_0|, 0)$ at all times and converges to 0 as $t \to \infty$ (because β is decreasing to 0 in the second argument). It turns out that this is exactly equivalent to global asymptotic stability (GAS) of the unforced system

$$\dot{x} = f(x, 0)$$

In this case one says that (82) is *0-GAS*.

A more standard definition of GAS is not in terms of class $\mathcal{KL}$ functions, but the two are equivalent; see [Khalil, Definition 4.4 and Lemma 4.5]. In particular, global exponential stability (GES) is defined as

$$|x(t)| \leq ce^{-\lambda t}|x_0|$$

and the right-hand side is an example of a class $\mathcal{KL}$ function.

However, GAS and GES are *internal stability* notions while ISS is an *external stability* notion. It says that in the presence of d, there's another term in the upper bound for the state, and this term depends on the size of the disturbance. The implications of ISS are:

- If d is bounded, then x is bounded.
- If $d \to 0$, then $x \to 0$.

The first one is obvious from the definition of ISS. The second one holds because we can always "restart" the system after d becomes small, and use (83). We already used this trick earlier to prove Fact 3 in Section 3.3, see an earlier homework problem.

Some examples:

The linear system

$$\dot{x} = Ax + Bd$$

is ISS if (and only if) A is a Hurwitz matrix. In other words, for linear systems internal stability (stability of $\dot{x} = Ax$) automatically implies external stability (ISS). We have

$$|x(t)| \leq ce^{-\lambda t}|x_0| + b \cdot \sup_{t_0 \leq s \leq t} |d(s)|$$

for suitable constants $b, c, \lambda > 0$. We actually already discussed this before (cf. Fact 2 in Section 3.3, also homework).

For nonlinear systems, it's no longer true that internal stability (GAS for $d \equiv 0$) implies ISS. We already saw this in the previous subsection. The system

$$\dot{x} = -x + xd$$

fails the first bullet item above (just set $d \equiv 2$). The system

$$\dot{x} = -x + x^2 d$$

fails both bullet items above—in fact, it fails miserably: not only does x not converge, not only is it unbounded, but it escapes to infinity in finite time. And all this for exponentially converging d, and despite the fact that for $d \equiv 0$ we have the nice stable linear system $\dot{x} = -x$.

It is possible to construct examples with less drastic behavior which fail the two bullet items above.

From the above discussion it is clear that a new theory is needed to study ISS. We will not pursue this in detail here, but we will mention a few basic results.

The ISS property admits the following Lyapunov-like equivalent characterization: the system (82) is ISS *if and only if* there exists a positive definite radially unbounded $\mathcal{C}^1$ function $V : \mathbb{R}^n \to \mathbb{R}$ such that for some class $\mathcal{K}_\infty$ function ρ we have

$$|x| \geq \rho(|d|) \quad \Longrightarrow \quad \frac{\partial V}{\partial x} \cdot f(x,d) < 0 \qquad \forall\, x \neq 0,\ \forall\, d \tag{84}$$

(to be read like this: for all $x \neq 0$ and all d, the implication holds). Such a function V is called an *ISS-Lyapunov function.*

Idea of the proof that the existence of an ISS-Lyapunov function implies ISS:

- As long as $|x| \geq \rho(|d|)$, we have $\dot{V} < 0$ and the system behaves like a usual GAS system, i.e., it converges towards the origin. During this period of time, we have an estimate of the form
$$|x(t)| \leq \beta(|x_0|, t)$$
- Assume that d is bounded (for unbounded d the ISS estimate gives no finite bound and so there's nothing to prove). Once we enter a level set of V superscribed around the ball of radius
$$\rho\big(\sup_{t_0 \leq s < \infty} |d(s)|\big)$$
we may no longer have $\dot{V} < 0$, but we know that we cannot exit this level set because outside it, $\dot{V} < 0$ and we are pushed back in. So, from this time onward, we satisfy the bound of the form
$$|x(t)| \leq \gamma\big(\sup_{t_0 \leq s < \infty} |d(s)|\big)$$
Here γ is obtained from ρ and V (geometrically, the ball of radius γ is superscribed around the level set of V which is in turn superscribed around the ball of radius ρ).
- Combining the two bounds—the one before we enter the level set of V and the one after—we obtain ISS.

The proof of the converse implication is much more difficult.

Begin optional material

As we said, 0-GAS (GAS under zero input) does not imply ISS. However, 0-GAS does imply ISS-like properties for sufficiently small inputs.

To see why, let V be a Lyapunov function for $\dot{x} = f(x,0)$, and write

$$\frac{\partial V}{\partial x} \cdot f(x,d) = \frac{\partial V}{\partial x} \cdot f(x,0) + \frac{\partial V}{\partial x} \cdot (f(x,d) - f(x,0)) \leq -W(x) + \left|\frac{\partial V}{\partial x}\right| L|d|$$

where W is positive definite and L is the Lipschitz constant for f on a region containing the initial condition x_0. It is now easy to see that for d small enough, the negative term dominates. We can now establish the ISS estimate as in the previous argument.

So, the key thing about ISS is that the estimate (83) holds for *arbitrary* inputs, no matter how large. This is a much stronger property and it doesn't follow from 0-GAS.
End optional material

So, to check ISS, we need to look for a *pair* (V, ρ) satisfying (84). We'll see examples of how to do this very soon.

Suppose now that we have a system with both disturbance inputs and control inputs:

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

Then a natural problem to consider is to design a feedback law $u = k(x)$ such that the closed-loop system is ISS with respect to the disturbance d. Such a control law is called *input-to-state stabilizing* (it attenuates the disturbance in the ISS sense).

Combining the notion of control Lyapunov function (CLF) discussed in Section 4.1 with the above notion of ISS-Lyapunov function, we arrive at the following definition of *ISS control Lyapunov function (ISS-CLF)*:

$$|x| \geq \rho(|d|) \quad \Longrightarrow \quad \inf_u \left\{ \frac{\partial V}{\partial x} \cdot f(x,d,u) \right\} < 0 \qquad \forall\, x \neq 0,\ \forall\, d$$

where $\rho \in \mathcal{K}_\infty$.

Given an ISS-CLF, we want to have a systematic procedure—in fact, a *universal formula*—for designing an input-to-state stabilizing controller, with V serving as an ISS-Lyapunov function for the closed loop.

We remember from the earlier case of no disturbances that to get this, we need to impose an affine structure on the system. Namely, let us assume that the right-hand side is affine in both[9] u and d:

$$\dot{x} = f(x) + L(x)d + G(x)u \tag{85}$$

Then V is an ISS-CLF if

$$|x| \geq \rho(|d|) \quad \Longrightarrow \quad \inf_u \left\{ \frac{\partial V}{\partial x} \cdot f(x) + \frac{\partial V}{\partial x} \cdot L(x)d + \frac{\partial V}{\partial x} \cdot G(x)u \right\} < 0 \qquad \forall\, x \neq 0,\ \forall\, d$$

[9] Affine dependence on d is not necessary, but without it the construction is more complicated [L-Sontag-Wang]. Affine dependence on u is essential.

It is still not clear how to apply Sontag's universal formula to this. First, the conditions are stated differently than before (in terms of the implication). Second, the expression inside the inf involves d, while we want the controller to be independent of d (which is usually not measured[10]).

(Recall the CLF setting: want to have $\inf_u\{a(x)+b(x)\cdot u\}<0$. Cannot define a in the obvious way because it'll depend on d.)

The trick is to realize that

$$\frac{\partial V}{\partial x}\cdot L(x)d \le \left|L^T(x)\frac{\partial V}{\partial x}\right| |d|$$

and that under the condition $|x|\ge\rho(|d|)$, the worst-case value of the above expression is

$$\left|L^T(x)\frac{\partial V}{\partial x}\right| \left|\rho^{-1}(|x|)\right|$$

This is well-defined because ρ, being a class $\mathcal{K}_\infty$ function, is invertible on $[0,\infty)$ (since it's strictly increasing from 0 to ∞). Also, there does in fact exist an admissible disturbance for which the above upper bound is achieved: just align d with the vector $L^T(x)\dfrac{\partial V}{\partial x}$.

We conclude that the following is an equivalent characterization of an ISS-CLF for the affine system (85):

$$\inf_u\left\{\frac{\partial V}{\partial x}\cdot f(x)+\left|L^T(x)\frac{\partial V}{\partial x}\right|\left|\rho^{-1}(|x|)\right|+\frac{\partial V}{\partial x}\cdot G(x)u\right\}<0 \qquad \forall\, x\ne 0$$

This looks more familiar. Defining

$$a(x):=\frac{\partial V}{\partial x}\cdot f(x)+\left|L^T(x)\frac{\partial V}{\partial x}\right|\left|\rho^{-1}(|x|)\right|, \qquad b(x):=G^T(x)\frac{\partial V}{\partial x}$$

we can rewrite the property of being an ISS-CLF as

$$\inf_u\{a(x)+b(x)\cdot u\}<0 \qquad \forall\, x\ne 0$$

or, equivalently, as

$$|b(x)|=0 \quad\Longrightarrow\quad a(x)<0$$

for all $x\ne 0$, which is exactly (36). And neither a nor b depends on d.

Now, the desired input-to-state stabilizing feedback law $u=k(x)$ is given by the universal formula (37). (If one wants this feedback law to be smooth, then one needs to replace the second term in $a(x)$, which is in general just continuous, by a smooth approximation.)

Claim: The closed-loop system is ISS.

[10] One exception arises in switching adaptive control, to be discussed later, where the disturbance corresponds to the output estimation error which is available for control.

Indeed, the derivative of V along the closed-loop system is

$$\dot{V} = \frac{\partial V}{\partial x} \cdot (f(x) + L(x)d + G(x)k(x))$$

and we have

$$|x| \geq \rho(|d|)$$
$$\Downarrow$$
$$\dot{V} \leq \frac{\partial V}{\partial x} \cdot f(x) + \left| L^T(x) \frac{\partial V}{\partial x} \right| \left| \rho^{-1}(|x|) \right| + \frac{\partial V}{\partial x} \cdot G(x)k(x) = a(x) + b(x) \cdot k(x) < 0$$

for all $x \neq 0$, where the last inequality (< 0) is guaranteed by the universal formula. Thus V is an ISS-Lyapunov function for the closed-loop system, which implies ISS by a result given earlier.

7.2.1 ISS backstepping

Now we'd like to design an ISS controller for the system

$$\begin{aligned} \dot{x} &= f(x) + L(x)d + G(x)\xi \\ \dot{\xi} &= u \end{aligned}$$

Before we do this, we need one more fact about ISS-Lyapunov functions. We defined ISS-Lyapunov functions via

$$|x| \geq \rho(|d|) \quad \Longrightarrow \quad \frac{\partial V}{\partial x} \cdot f(x, d) < 0 \qquad \forall\, x \neq 0,\ \forall\, d$$

where $\rho \in \mathcal{K}_\infty$. We could also rewrite this more precisely as

$$|x| \geq \rho(|d|) \quad \Longrightarrow \quad \frac{\partial V}{\partial x} \cdot f(x, d) \leq -\alpha(|x|) \qquad \forall\, x, d \tag{86}$$

where $\alpha \in \mathcal{K}_\infty$ (the fact that we can take α to be of class $\mathcal{K}_\infty$, and not just positive definite, is not obvious, but this can be always be achieved by modifying V). Another equivalent characterization of ISS-Lyapunov functions is as follows:

$$\frac{\partial V}{\partial x} \cdot f(x, d) \leq -\alpha(|x|) + \chi(|d|) \qquad \forall\, x, d \tag{87}$$

where $\alpha, \chi \in \mathcal{K}_\infty$.

Begin optional material ______________________________

Proof of the equivalence (sketch). It is not hard to obtain χ from ρ and vice versa:

Suppose that (87) holds. Rewrite it as

$$\frac{\partial V}{\partial x} \cdot f(x, d) \leq -\frac{1}{2}\alpha(|x|) - \frac{1}{2}\alpha(|x|) + \chi(|d|)$$

from which we see that

$$|x| \geq \alpha^{-1}(2\chi(|d|)) \quad \Longrightarrow \quad \frac{\partial V}{\partial x} \cdot f(x,d) \leq -\frac{1}{2}\alpha(|x|)$$

(In this case we get not α but $\frac{1}{2}\alpha$ on the right-hand side, but the constant $\frac{1}{2}$ is arbitrary and it could be arbitrarily close to 1.)

Conversely, suppose that (86) holds. We only need to show (87) when

$$|x| \leq \rho(|d|)$$

which can be done because

$$\max_{|x| \leq \rho(|d|)} \left\{ \frac{\partial V}{\partial x} \cdot f(x,d) + \alpha(|x|) \right\}$$

is a function of d only and hence

$$\chi(r) := \max_{|d| \leq r,\, |x| \leq \rho(r)} \left\{ \frac{\partial V}{\partial x} \cdot f(x,d) + \alpha(|x|) \right\}$$

is well defined and gives (87). □

End optional material

Accordingly, we could have equivalently defined ISS-CLFs via

$$\inf_u \left\{ \frac{\partial V}{\partial x} \cdot f(x,d,u) \right\} \leq -\alpha(|x|) + \chi(|d|) \qquad \forall\, x, d$$

For applying the universal formula, this would not have been as convenient (the term $\chi(|d|)$ outside the inf complicates things, it needs to be moved inside first). But for backstepping, this alternative formulation will be more convenient. More precisely, we will see that it's not useful for the initialization step of backstepping but useful for the recursion.

Lecture 17 (Mon Nov 2)

We're now ready for ISS backstepping. Suppose that we're given an ISS-CLF $V_0(x)$ and a corresponding control law $k_0(x)$, smooth and satisfying $k_0(0) = 0$, for which we have

$$\frac{\partial V_0}{\partial x} \cdot (f(x) + L(x)d + G(x)k_0(x)) \leq -\alpha_0(|x|) + \chi_0(|d|)$$

We claim that the "augmented" function defined in the usual way:

$$V_1(x,\xi) := V_0(x) + \frac{1}{2}|\xi - k_0(x)|^2$$

is an ISS-CLF for the augmented system, and we'll show this by explicitly constructing the new input-to-state stabilizing control law $k_1(x,\xi)$.

The derivative of V_1 is given by (k_0' stands for the Jacobian matrix $\frac{\partial k_0}{\partial x}$)

$$\begin{aligned}\dot V_1 &= \underbrace{\frac{\partial V_0}{\partial x} \cdot (f + Ld + Gk_0)}_{\text{“old” } \dot V_0} + \frac{\partial V_0}{\partial x} \cdot G(\xi - k_0) + (\xi - k_0)^T(u - k_0'f - k_0'Ld - k_0'G\xi) \\ &\leq -\alpha_0(|x|) + \chi_0(|d|) + (\xi - k_0)^T\left(u - k_0'f - k_0'Ld - k_0'G\xi + G^T\frac{\partial V_0}{\partial x}\right)\end{aligned}$$

We can cancel all terms inside the parentheses, except $k_0'Ld$. But we can dominate this term using square completion: define

$$k_1(x,\xi) := -(\xi - k_0) + k_0'f + k_0'G\xi - G^T\frac{\partial V_0}{\partial x} - (k_0'L)(k_0'L)^T(\xi - k_0)$$

Then we get

$$\begin{aligned}\dot V_1 &\leq -\alpha_0(|x|) - |\xi - k_0(x)|^2 + \chi_0(|d|) - (\xi - k_0)^T(k_0'L)(k_0'L)^T(\xi - k_0) - (\xi - k_0)^T k_0'Ld \\ &= -\alpha_0(|x|) - |\xi - k_0(x)|^2 + \chi_0(|d|) \\ &\qquad \underbrace{-(\xi - k_0)^T(k_0'L)(k_0'L)^T(\xi - k_0) - (\xi - k_0)^T k_0'Ld - \frac{1}{4}d^Td}_{= -\left|(k_0'L)^T(\xi - k_0) + \frac{1}{2}d\right|^2} + \frac{1}{4}d^Td \\ &\leq -\alpha_0(|x|) - |\xi - k_0(x)|^2 + \chi_0(|d|) + \frac{1}{4}d^Td \leq -\alpha_1\left(\left|\begin{pmatrix} x \\ \xi \end{pmatrix}\right|\right) + \chi_1(|d|)\end{aligned}$$

where

$$\chi_1(|d|) := \chi_0(|d|) + \frac{1}{4}|d|^2$$

and $\alpha_1 \in \mathcal{K}_\infty$ can be suitably defined since $\alpha_0(|x|) + |\xi - k_0(x)|^2$ is positive definite and radially unbounded as a function of (x, ξ).

As before, we can apply the above backstepping procedure recursively to handle a chain of integrators.

$\longrightarrow$ For *initializing* the backstepping procedure, it is useful to know that any *scalar* affine system

$$\dot{x} = f(x) + \ell(x)d + g(x)u$$

with $g(x) \neq 0$ is input-to-state stabilized by the feedback

$$u = \frac{1}{g(x)}\big(- f(x) - x - |\ell(x)|x\big)$$

Indeed, the closed-loop system is

$$\dot{x} = -x + \ell(x)d - |\ell(x)|x$$

and for

$$V(x) = \frac{x^2}{2}$$

we have

$$\begin{aligned} \dot{V} = -x^2 + \ell(x)dx - |\ell(x)|x^2 \le -x^2 + |\ell(x)||d||x| - |\ell(x)|x^2 &= -x^2 - |\ell(x)||x|(|x| - |d|) \\ &\le -x^2 \quad \text{if} \quad |x| \ge |d| \end{aligned}$$

hence V is an ISS-Lyapunov function.

(We also know that in general, the assumption $g(x) \neq 0$ is not necessary for V to be an ISS-CLF.)

As before, this can be generalized to strict feedback systems (cf. earlier homework).

7.3 Adaptive ISS controller design

Reference: [KKK book, Section 5.2]

We now go back to adaptive control and continue the discussion we started in Section 7.1. There, the disturbance was the parameter estimation error. So, we'll now try to adopt the above general theory of ISS controller design to this more concrete scenario.

$\longrightarrow$ We will confine ourselves to a specific example. The general formulas are quite similar to those we derived for the non-adaptive case. The example will be enough to bring out the differences between the non-adaptive and adaptive cases.

Consider again the scalar plant from Example 9:

$$\dot{x} = -\theta x^2 + u$$

We saw that if we use the controller

$$u = -x + \hat{\theta} x^2$$

then the closed-loop system

$$\dot{x} = -x + \tilde{\theta}x^2$$

is *not* ISS with respect to the parameter estimation error $\tilde{\theta}$. As a result, if $\tilde{\theta}$ is bounded or even converges to 0, convergence or even boundedness of x is not guaranteed. We recognize this as lack of ISS property with respect to $\tilde{\theta}$.

Can we find a different controller that does provide ISS with respect to $\tilde{\theta}$?

We need to add some term that will dominate x^2 for large x. With this intuition, one obvious choice is

$$u = -x + \hat{\theta}x^2 - x^3$$

The closed-loop system is

$$\dot{x} = -x + \tilde{\theta}x^2 - x^3$$

It is quite clear that solutions will no longer escape to ∞. But we need to prove ISS.

Since the system is scalar, we take the candidate ISS-Lyapunov function

$$V(x) = \frac{x^2}{2}$$

for which

$$\dot{V} = -x^2 + \tilde{\theta}x^3 - x^4 \leq -x^2 + |\tilde{\theta}||x^3| - x^4 = -x^2 - |x^3|(|x| - |\tilde{\theta}|)$$

and we get

$$|x| \geq |\tilde{\theta}| \quad \Longrightarrow \quad \dot{V} = -x^2 < 0 \qquad \forall\, x \neq 0$$

This by definition shows that V is an ISS-Lyapunov function, hence the system is indeed ISS.

Remark 1 We're viewing $\tilde{\theta}$ as an external input. However, once the tuning law which generates $\hat{\theta}$ is in place, it is more accurate to think of $\tilde{\theta}$ as an *output* of the overall adaptive system, and not as an input. This is because $\tilde{\theta}$ is determined by θ and $\hat{\theta}$, which are the internal system parameters. Then, a more accurate term is *output-to-state stability (OSS)*, which is a form of detectability (small output implies small state, etc.) This doesn't change much as far as the theory goes, but it's helpful for achieving conceptual clarity. Detectability is very important in the adaptive control context, we'll see more on this in Section 8.3 and Section 9. □

7.3.1 Adaptive ISS backstepping

To go beyond scalar systems, we can again apply backstepping. More precisely, we need to develop an adaptive version of ISS backstepping discussed above. As in the case of usual (not ISS) backstepping, discussed in Section 5, the basic idea will carry over but there will be some differences.

Let us see how it works by adding an integrator to the previous system:

$$\dot{x} = -\theta x^2 + \xi$$
$$\dot{\xi} = u$$

We already did the initial step, which we reproduce below with appropriate notation:

$$u = -x + \hat{\theta}x^2 - x^3 =: k_0(x, \hat{\theta})$$

$$V_0(x) = \frac{x^2}{2}$$

$$\frac{\partial V_0}{\partial x}(-\theta x^2 + k_0) = x(-x + \tilde{\theta}x^2 - x^3) \le -x^2 - |x^3|(|x| - |\tilde{\theta}|) \le -x^2 \quad \text{if} \quad |x| \ge |\tilde{\theta}|$$

But we know that this "gain-margin" characterization of ISS-Lyapunov function is not the right one to use for backstepping. We want one of the form (87). This is easy to get: when

$$|x| \le |\tilde{\theta}|$$

we can rewrite the above expression for $\dot{V}_0$ as

$$-x^2 + |x^3|(|\tilde{\theta}| - |x|) \le -x^2 + |\tilde{\theta}^3||\tilde{\theta}| = -x^2 + \tilde{\theta}^4$$

Combining the two cases, we arrive at

$$\frac{\partial V_0}{\partial x}(-\theta x^2 + k_0) \le -x^2 + \tilde{\theta}^4$$

Now, the usual candidate ISS-CLF for the augmented system:

$$V_1(x, \hat{\theta}) := V_0(x) + \frac{1}{2}|\xi - k_0(x, \hat{\theta})|^2 = \frac{x^2}{2} + \frac{1}{2}(\xi + x - \hat{\theta}x^2 + x^3)^2$$

$\longrightarrow$ Note: this Lyapunov function depends on $\hat{\theta}$, because so does the control law k_0. As we'll see in a moment, this introduces additional complications compared to the non-adaptive case.

$$\begin{aligned}\dot{V}_1 &= \frac{\partial V_0}{\partial x}(-\theta x^2 + \xi) + (\xi + x - \hat{\theta}x^2 + x^3)\big(u + (1 - 2\hat{\theta}x + 3x^2)(-\theta x^2 + \xi) - x^2\dot{\hat{\theta}}\big) \\ &= \underbrace{x(-x + \tilde{\theta}x^2 - x^3)}_{\le -x^2 + \tilde{\theta}^4} + (\xi + x - \hat{\theta}x^2 + x^3)\big(u + x + (1 - 2\hat{\theta}x + 3x^2)(-\hat{\theta}x^2 + \xi) \\ &\qquad\qquad + (1 - 2\hat{\theta}x + 3x^2)x^2\tilde{\theta} - x^2\dot{\hat{\theta}}\big)\end{aligned}$$

(As earlier on page 46, we used $\theta = \hat{\theta} - \tilde{\theta}$ to split the θ-dependent term.) We know that we can cancel the terms that come after u on the second line of the above formula, and add damping:

$$k_1(x, \hat{\theta}, \xi) := -x - (1 - 2\hat{\theta}x + 3x^2)(-\hat{\theta}x^2 + \xi) - (\xi + x - \hat{\theta}x^2 + x^3)\ldots \quad \text{(not yet complete)}$$

What should we do about the term multiplied by $\tilde{\theta}$?

This is the disturbance-dependent term, which in the non-adaptive case was $k_0' Ld$. As we did there, we can dominate it by square completion:

$$\begin{aligned} k_1(x, \hat{\theta}, \xi) := &- x - (1 - 2\hat{\theta}x + 3x^2)(-\hat{\theta}x^2 + \xi) - (\xi + x - \hat{\theta}x^2 + x^3) \\ &- (\xi + x - \hat{\theta}x^2 + x^3)(1 - 2\hat{\theta}x + 3x^2)^2 x^4 \ldots \text{ (not yet complete)} \end{aligned}$$

and this will contribute the term

$$\frac{1}{4}\tilde{\theta}^2$$

to the bound for $\dot{V}_1$ (because we have to subtract and add this term to complete the square).

What should we do about the term $-x^2\dot{\hat{\theta}}$?

Actually, there's nothing we can really do about this term! In our previous adaptive backstepping design (Section 5.2), we allowed u to depend on this term, i.e., we allowed u to depend on $\dot{\hat{\theta}}$ and thus be coupled to the tuning law. But here we are aiming for *modular design*, i.e., we want the control law and tuning law designs to be decoupled.

$\longrightarrow$ To achieve this, we need to *reformulate our ISS objective.* In addition to the parameter estimation error $\tilde{\theta}$, we will view its derivative $\dot{\tilde{\theta}} = \dot{\hat{\theta}}$ as an *additional input* and ask for ISS with respect to the vector input

$$\begin{pmatrix} \tilde{\theta} \\ \dot{\tilde{\theta}} \end{pmatrix}$$

Thus the ISS notion we're now asking for is weaker than ISS with respect to $\tilde{\theta}$ only. (Just recall the definition of ISS; clearly, ISS with respect to a part of the input vector implies ISS with respect to the whole input vector, but not vice versa.)

Now that we're treating $\dot{\hat{\theta}}$ as another input, we can handle the last term in exactly the same way as the previous term, i.e., dominate it using square completion. The complete formula for the control law stands as

$$\begin{aligned} k_1(x, \hat{\theta}, \xi) := &- x - (1 - 2\hat{\theta}x + 3x^2)(-\hat{\theta}x^2 + \xi) - (\xi + x - \hat{\theta}x^2 + x^3) \\ &- (\xi + x - \hat{\theta}x^2 + x^3)(1 - 2\hat{\theta}x + 3x^2)^2 x^4 - (\xi + x - \hat{\theta}x^2 + x^3)x^4 \end{aligned}$$

This will contribute the term

$$\frac{1}{4}\dot{\tilde{\theta}}^2$$

to the bound for $\dot{V}_1$, which in the end will be

$$\dot{V}_1 \le -x^2 - (\xi + x - \hat{\theta}x^2 + x^3)^2 + \tilde{\theta}^4 + \frac{1}{4}\tilde{\theta}^2 + \frac{1}{4}\dot{\tilde{\theta}}^2 \le -\alpha_1\left(\left|\begin{pmatrix} x \\ \xi \end{pmatrix}\right|\right) + \left|\begin{pmatrix} \tilde{\theta} \\ \dot{\tilde{\theta}} \end{pmatrix}\right|^4 + \frac{1}{4}\left|\begin{pmatrix} \tilde{\theta} \\ \dot{\tilde{\theta}} \end{pmatrix}\right|^2$$

for a suitable $\alpha_1 \in \mathcal{K}_\infty$, and ISS is established.

Since the control law k_1 doesn't depend on $\dot{\hat{\theta}}$, we can see that the procedure can be repeated if another integrator is added. So, we have a recursive procedure which can handle chains of integrators. And this procedure is cleaner than what we had in Section 5.2 using tuning functions.

Lecture 18 (Wed Nov 4)

7.3.2 Modular design

Now, suppose that our plant is

$$\dot{x} = f(x) + L(x)\theta + G(x)u$$

and we have an adaptive control law which guarantees ISS of the closed-loop system with respect to $(\tilde{\theta}, \dot{\hat{\theta}})$. Then, we know that the following is true:

- If $\tilde{\theta}$ and $\dot{\hat{\theta}}$ are bounded, then x is bounded
- If $\tilde{\theta}$ and $\dot{\hat{\theta}}$ converge to 0, then x converges to 0

Thus to guarantee the desired properties of the state x, we need to find an adaptive law which ensures the corresponding properties of the parameter estimation error $\tilde{\theta}$. This is what we meant by "modular design": here are the design objectives for the controller and the estimation law, and we can go ahead and design them separately. It doesn't matter exactly how we'll design them; as long as the above properties are satisfied, the problem will be solved. (Go back again to Section 5.2: do you see that we didn't have modularity there?)

We're not going to discuss the design of parameter estimation laws for nonlinear systems. But for linear systems, we studied this in detail in Chapter 6. In particular, we saw how to get boundedness of $\tilde{\theta}$ and $\dot{\hat{\theta}}$, as well as (in some cases, at least without normalization) convergence of $\dot{\hat{\theta}}$ to 0. We also had such properties in indirect MRAC. So, while we may not always know how to achieve these properties for general systems, they are not surprising and we're comfortable with them.

We're less comfortable, however, with the requirement

$$\tilde{\theta} \to 0$$

We know that we do not have this unless we have PE conditions (and we don't want to impose PE because it typically conflicts with the control design objective). In this context, it is helpful to know that ISS with respect to $(\tilde{\theta}, \dot{\hat{\theta}})$ also implies the following.

- If $\tilde{\theta}$ and $\dot{\hat{\theta}}$ are bounded and $L(x)\tilde{\theta}$ and $\dot{\hat{\theta}}$ converge to 0, then x converges to 0

In other words, $\tilde{\theta} \to 0$ is replaced by the less demanding condition that $L(x)\tilde{\theta} \to 0$. The proof of this fact relies on the result mentioned in the optional material in the ISS subsection. For more details, see [KKK book, Lemma 5.3 on p. 193].

Just to get an idea how we might get this latter convergence property, consider the estimator

$$\dot{\hat{x}} = A_m(\hat{x} - x) + f(x) + L(x)\hat{\theta} + G(x)u$$

where A_m is a Hurwitz matrix. The state estimation error

$$e = \hat{x} - x$$

satisfies

$$\dot{e} = Ae + L(x)\tilde{\theta}$$

Suppose that we were able to show that all signals are bounded and $e \to 0$ (again, for linear plants we were often able to get such properties in Chapter 6). Then $\dot{e}$, $\ddot{e}$ are bounded ($\dot{\tilde{\theta}}$ is bounded because it's some function of the other signals which are bounded). Also,

$$\int_0^\infty \dot{e}(s)ds = \lim_{t\to\infty} \int_0^t \dot{e}(s)ds = \lim_{t\to\infty} e(t) - e(0) = -e(0)$$

is bounded. Applying Barbalat's lemma (this means applying it *twice*, because one first already applies it once to show that $e \to 0$), we conclude that

$$\dot{e} \to 0$$

In view of the expression for $\dot{e}$ and the fact that $e \to 0$, we have

$$L(x)\tilde{\theta} \to 0$$

as needed.

For details on designing parameter estimation schemes providing the above properties, read [KKK book, Chapters 5 and 6].

8 Stability of slowly time-varying systems

We already had to deal with stability of linear time-varying systems (using observability). Need to understand this subject better.

8.1 Stability

Suppose that we're given an LTV system

$$\dot{x} = A(t)x \tag{88}$$

and we want to know when it is uniformly asymptotically stable (this stability is then automatically global and exponential by linearity).

Is it enough to assume that $A(t)$ is a Hurwitz matrix for each *fixed* t?

No! (Not even if eigenvalues of $A(t)$ are bounded away from the imaginary axis.)

Can we come up with a counterexample?

To understand this, the easiest case to consider is when $A(t)$ is a piecewise-constant function of time, i.e., it *switches* between several fixed matrices. Then instead of a more usual time-varying system we have a *switched system*. The two are closely related, and later we'll study switched systems more explicitly when we discuss switching adaptive control.

Counterexample:

Suppose that we are switching between two systems in the plane. Suppose that the two individual systems are asymptotically stable, with trajectories as shown in the first figure (the solid curve and the dotted curve).

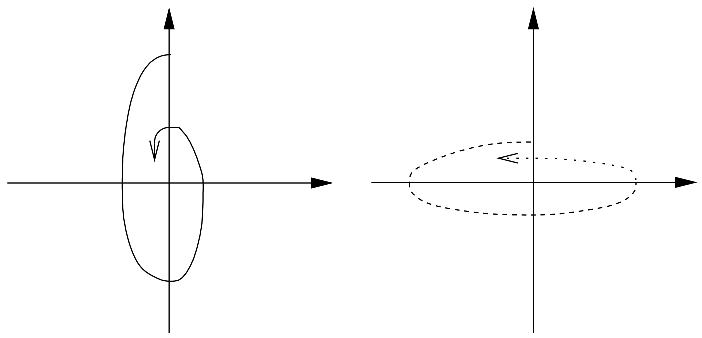

For different choices of the switching sequence, the switched system might be asymptotically stable or unstable; these two possibilities are shown in the second figure.

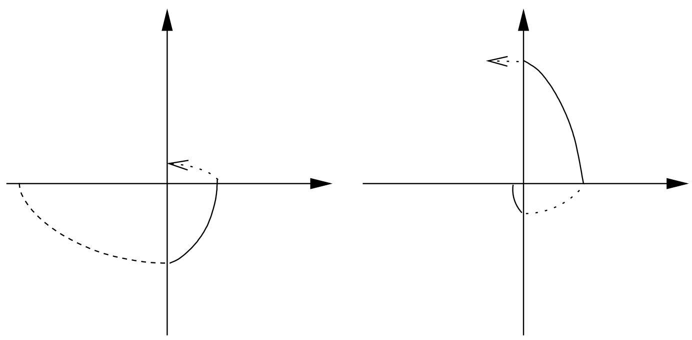

From this example, the instability mechanism is also quite clear: even though each system is stable, we catch it at the "peak" of its transient and switch to the other system, without giving it a chance to decay.

It is easy to modify this example to make $A(t)$ continuously time-varying (just imagine a family of systems homotopically connecting these two, and sliding fast through them would mimic the effect of a switch).

There are many results dealing with stability of switched and time-varying systems. In fact, stability of the constituent systems for frozen t is neither sufficient nor necessary (we could turn things around in the above example and have unstable systems but with switching having a stabilizing effect).

The direction we pursue here is suggested by the above example, and is also the one that will be relevant for adaptive control. As the section title suggests, we will restrict attention to time-variations that are sufficiently slow (assuming still that $A(t)$ is stable for each t). In the above example, it is clear that if we waited longer, we'd be OK. Here's a general result that characterizes such slow variations by imposing suitable conditions on the derivative $\dot{A}(\cdot)$.

Theorem 10 (Ioannou-Sun, Theorem 3.4.11) *Consider the LTV system* (88) *and assume that:*

- *$A(t)$ is Hurwitz for each fixed t, and there exist constants $c, \lambda_0 > 0$ such that for all t and s we have*[11]
$$\left\| e^{A(t)s} \right\| \leq ce^{-\lambda_0 s} \tag{89}$$
- *$A(\cdot)$ is C^1 and uniformly bounded: there exists an $L > 0$ such that $\|A(t)\| \leq L \; \forall \, t$.*
- *For all t and all T we have either*
  a) $\int_t^{t+T} \|\dot{A}(s)\| ds \leq \mu T + \alpha$
  or
  b) $\int_t^{t+T} \|\dot{A}(s)\|^2 ds \leq \mu T + \alpha$
  where $\mu, \alpha > 0$ and μ is sufficiently small.

Then the system (88) *is exponentially stable.*

The condition (89) means that all matrices $A(t)$ have a common stability margin λ_0 and a common overshoot constant c.

$\longrightarrow$ Note that $A(\cdot)$ satisfies one of the hypotheses a), b) if $\|\dot{A}(\cdot)\|$ is in L_∞ and sufficiently small, or if it is in L_1 or L_2.

When $\dot{A}(\cdot)$ satisfies a), it is sometimes called *nondestabilizing with growth rate μ*. If $\|\dot{A}(\cdot)\|$ is in L_1 then the growth rate is 0.

$\longrightarrow$ In the proof we'll get a formula for μ (an upper bound). However, we view this result more as a qualitative one (which is how it is stated).

Proof of Theorem 10.

Let's prove a).

For each fixed t, let $P(t)$ be the symmetric positive-definite solution of the Lyapunov equation

$$P(t)A(t) + A^T(t)P(t) = -I \tag{90}$$

[11] Here $\|\cdot\|$ stands for induced matrix norm corresponding to Euclidean norm.

$\longrightarrow$ This is pointwise in t ("system snapshot"). Don't confuse this with the Lyapunov equation for general LTV systems, which would be $\dot{P}(t) + P(t)A(t) + A^T(t)P(t) = -I$. Of course $\dot{P}(t)$ will eventually arise here too, but we'll use the slow-varying conditions to bound it. Basically, we first treat the system as if it were LTI and then use perturbation arguments.

Candidate Lyapunov function:

$$V(t, x) := x^T(t)P(t)x(t)$$

Its derivative:

$$\dot{V} = -|x|^2 + x^T \dot{P} x$$

We need to show some properties for $P(\cdot)$ and $\dot{P}(\cdot)$.

First, it is well known (ECE 515) that since each $P(t)$ is the solution of the corresponding Lyapunov equation (90), it's given by the formula

$$P(t) = \int_0^\infty e^{A^T(t)s} e^{A(t)s} ds \tag{91}$$

so from (89) we have

$$\|P(t)\| \le \int_0^\infty c^2 e^{-2\lambda_0 s} ds =: \beta_2$$

so $P(t)$ is uniformly bounded (L_∞).

Second, since $A(t)$ is uniformly bounded and is Hurwitz for each fixed t, we also have a lower bound

$$\left\|e^{A(t)s}\right\| \ge \bar{c} e^{-\bar{\lambda}_0 s} \qquad \forall\, t, s$$

This just says that since eigenvalues of $A(t)$ cannot be arbitrarily negative, there is a lower bound on how fast the solutions can decay for the LTI systems corresponding to $A(t)$ for all frozen t. Then, using (91) again, we similarly get

$$\|P(t)\| \ge \beta_1 > 0$$

so $P(t)$ is uniformly bounded away from 0. (Cf. [Khalil, p. 159 or 371].)

Now we differentiate (90):

$$\dot{P}(t)A(t) + A^T(t)\dot{P}(t) = -P(t)\dot{A}(t) - \dot{A}^T(t)P(t) =: -Q(t)$$

which implies

$$\dot{P}(t) = \int_0^\infty e^{A^T(t)s} Q(t) e^{A(t)s} ds \quad \Longrightarrow \quad \|\dot{P}(t)\| \le \|Q(t)\| \int_0^\infty c^2 e^{-2\lambda_0 s} ds \le \beta_2 \|Q(t)\|$$

From definition of Q,

$$\|Q(t)\| \le 2\|P(t)\|\|\dot{A}(t)\| \le 2\beta_2 \|\dot{A}(t)\|$$

Combining the two:

$$\|\dot{P}(t)\| \leq 2\beta_2^2 \|\dot{A}(t)\|$$

Plugging this into the earlier formula for $\dot{V}$:

$$\dot{V} \leq -|x|^2 + 2\beta_2^2 \|\dot{A}(t)\| |x|^2$$

From the definition of V and the above bounds we have the following:

$$\beta_1 |x|^2 \leq V(t,x) \leq \beta_2 |x|^2$$

This gives

$$\dot{V} \leq -\beta_2^{-1} V + 2\beta_2^2 \beta_1^{-1} \|\dot{A}(t)\| V = -\left(\beta_2^{-1} - 2\beta_2^2\beta_1^{-1}\|\dot{A}(t)\|\right) V$$

By standard comparison principle,

$$V(t) \leq e^{-\int_{t_0}^{t} (\beta_2^{-1} - 2\beta_2^2\beta_1^{-1}\|\dot{A}(s)\|)ds} V(t_0) \underset{\text{by a)}}{\leq} e^{2\beta_2^2\beta_1^{-1}\alpha} e^{-(\beta_2^{-1} - 2\beta_2^2\beta_1^{-1}\mu)(t-t_0)} V(t_0)$$

Therefore, $V(t)$ decays to 0 exponentially if

$$\mu < \frac{\beta_1}{2\beta_2^3}$$

To prove b), one uses Cauchy-Schwartz inequality to write

$$\int_{t_0}^{t} \|\dot{A}(s)\| ds \leq \sqrt{\int_{t_0}^{t} \|\dot{A}(s)\|^2 ds} \cdot \sqrt{t-t_0} \leq \sqrt{\mu(t-t_0)^2 + \alpha(t-t_0)} \leq \sqrt{\mu}(t-t_0) + \sqrt{\alpha}\sqrt{t-t_0}$$

and notes that the term $\sqrt{t-t_0}$ inside the exponential will be dominated by $-\beta_2^{-1}(t-t_0)$. See [Ioannou-Sun] for details. □

8.2 Application to adaptive stabilization

In the context of adaptive control, time-variation comes from on-line adjustments of the parameter estimates $\hat{\theta}(t)$, and slow time-variation in the sense of Theorem 10 corresponds to slow speed of adaptation.

Recall that adaptive laws we studied earlier (see, e.g., Theorem 5 in Section 6.3) provide slow adaptation in the sense that $\dot{\hat{\theta}}(t)$ is in L_∞ and L_2. This makes Theorem 10 applicable.

Lecture 19 (Mon Nov 9)

In adaptive control, our basic idea has been to design the control law based on certainty equivalence and combine it with an adaptive law. Then, for each frozen value of the parameter estimates the corresponding controller by construction provides some form of stability, and these parameter estimates don't change too fast.

We need to be a little more careful because for each fixed value of the parameter estimates, the controller stabilizes the corresponding plant model, but not the actual unknown plant. So, typically Theorem 10 is applied to a plant model and then a perturbation argument is used to show stability of the real closed-loop plant.

Consider the two-dimensional, single-input example

$$\dot{x} = \begin{pmatrix} a & 1 \\ 0 & 0 \end{pmatrix} x + \begin{pmatrix} b \\ b \end{pmatrix} u$$

where $b \neq 0$ and both a and b are unknown. (We don't assume that $a < 0$.)

The objective is to asymptotically stabilize this system ($x \to 0$) using state feedback.

$\longrightarrow$ A similar system arises in the context of adaptive tracking in [Ioannou-Sun], see examples in Sect. 7.4.3 and 7.4.4. (There is a preliminary step in converting the tracking problem to stabilization of the error system, which we skipped. The error system in [Ioannou-Sun] is similar to the above system, but it has more equations because there is also an observer part there, while here we assume that both components of x are available for control.)

To apply the certainty equivalence principle of control design, we first need to select a controller that stabilizes our system for the case when a, b are *known*.

It is easy to check that the system is controllable (hence stabilizable). We can stabilize it with a state feedback controller

$$v = -\begin{pmatrix} k_1 & k_2 \end{pmatrix} x = -k_1 x_1 - k_2 x_2 \tag{92}$$

where k_1 and k_2 are selected by any standard linear control design method discussed in ECE 515.

One option is to place closed-loop eigenvalues at some chosen locations in the left half of the complex plane (pole placement). The closed-loop matrix is

$$A_{\text{cl}} = \begin{pmatrix} a & 1 \\ 0 & 0 \end{pmatrix} + \begin{pmatrix} b \\ b \end{pmatrix} \begin{pmatrix} -k_1 & -k_2 \end{pmatrix} = \begin{pmatrix} a - bk_1 & 1 - bk_2 \\ -bk_1 & -bk_2 \end{pmatrix}$$

Choosing, e.g.,

$$k_1 = \frac{a+1}{b}, \qquad k_2 = \frac{1}{b}$$

we get

$$A_{\text{cl}} = \begin{pmatrix} -1 & 0 \\ -a-1 & -1 \end{pmatrix}$$

whose eigenvalues are $-1, -1$. (There is a systematic procedure for designing pole-placement feedback laws, via transforming to controllable canonical form, but here we can just find the gains by trial and error.)

Begin optional material ______

Another option is, instead of arbitrarily selecting closed-loop poles, to consider an LQR problem

$$J = \int_0^\infty (x^T(s)Qx(s) + Ru^2(s))ds \longrightarrow \min_u$$

where the matrix $Q = Q^T \geq 0$ and the scalar $R > 0$ are design parameters (they weight stability against control effort). The solution is

$$u = -\frac{1}{R}B^T Px$$

where $P = P^T > 0$ satisfies the algebraic Riccati equation

$$PA + A^T P - \frac{1}{R}PBB^T P + Q = 0$$

End optional material ______

For simplicity, we choose the first option and go with the above formulas for k_1 and k_2.

Now, back to the case when a and b are *unknown*, and we need to work with their estimates $\hat{a}$, $\hat{b}$ and the corresponding estimation errors

$$\tilde{a} := \hat{a} - a, \qquad \tilde{b} := \hat{b} - b$$

Rewrite the system as

$$\dot{x} = \begin{pmatrix} \hat{a} - \tilde{a} & 1 \\ 0 & 0 \end{pmatrix} x + \begin{pmatrix} \hat{b} - \tilde{b} \\ \hat{b} - \tilde{b} \end{pmatrix} u = \underbrace{\begin{pmatrix} \hat{a} & 1 \\ 0 & 0 \end{pmatrix} x + \begin{pmatrix} \hat{b} \\ \hat{b} \end{pmatrix} u}_{\text{design model}} \underbrace{- \begin{pmatrix} \tilde{a} \\ 0 \end{pmatrix} x_1 - \begin{pmatrix} \tilde{b} \\ \tilde{b} \end{pmatrix} u}_{\text{perturbations}}$$

The design model has the same form as the original system but instead of the unknown parameters we have their estimates which are available for control. We need a control law

$$u = -\hat{k}_1 x_1 - \hat{k}_2 x_2$$

such that the matrix

$$\widehat{A}_{\text{cl}} := \begin{pmatrix} \hat{a} & 1 \\ 0 & 0 \end{pmatrix} + \begin{pmatrix} \hat{b} \\ \hat{b} \end{pmatrix} \begin{pmatrix} -\hat{k}_1 & -\hat{k}_2 \end{pmatrix} = \begin{pmatrix} \hat{a} - \hat{b}\hat{k}_1 & 1 - \hat{b}\hat{k}_2 \\ -\hat{b}\hat{k}_1 & -\hat{b}\hat{k}_2 \end{pmatrix}$$

is Hurwitz. Repeating the above control design with $\hat{a}, \hat{b}$ in place of a, b, we get

$$\hat{k}_1 = \frac{\hat{a} + 1}{\hat{b}}, \qquad \hat{k}_2 = \frac{1}{\hat{b}}$$

and

$$\widehat{A}_{\text{cl}} = \begin{pmatrix} -1 & 0 \\ -\hat{a} - 1 & -1 \end{pmatrix} \tag{93}$$

whose eigenvalues are $-1, -1$.

Now, we need an adaptive law for updating the parameter estimates $\hat{a}$ and $\hat{b}$. We already developed parameter estimation schemes that will cover this example. We can use, e.g., the normalized gradient law of Section 6.3.2, whose properties are listed in Theorem 5.[12] That theorem guarantees, in particular:

$$\hat{a}, \hat{b} \in L_\infty, \qquad \dot{\hat{a}}, \dot{\hat{b}} \in L_2 \cap L_\infty \tag{94}$$

(only item (i) of Theorem 5 applies since we don't have persistency of excitation here).

$\longrightarrow$ The design of the control law and the parameter estimation law are completely independent of one another ("modular design").

Combining the two, we get the closed-loop system

$$\dot{x} = \widehat{A}_{\text{cl}}(t)x - \begin{pmatrix} \tilde{a} \\ 0 \end{pmatrix} x_1 - \begin{pmatrix} \tilde{b} \\ \tilde{b} \end{pmatrix} u \tag{95}$$

where we write the argument t in $\widehat{A}_{\text{cl}}$ to emphasize that it is time-varying because $\hat{a}$ evolves with time. (So does $\hat{b}$ but $\widehat{A}_{\text{cl}}$ doesn't depend on it.)

The analysis now proceeds in two steps:

Step 1 By construction, $\widehat{A}_{\text{cl}}(t)$ is Hurwitz for each t. Using (93) and (94), we see that

$$\left\| \dot{\widehat{A}}_{\text{cl}}(t) \right\| = \left\| \begin{pmatrix} 0 & 0 \\ -\dot{\hat{a}} & 0 \end{pmatrix} \right\| = |\dot{\hat{a}}| \in L_2$$

Hence, Theorem 10 applies and tells us that the part of (95) coming from the design model, i.e.,

$$\dot{x} = \widehat{A}_{\text{cl}}(t)x$$

is exponentially stable.

$\longrightarrow$ **I forgot to mention this in class:** What if $\hat{b}(t) = 0$ at some time t? This is *loss of stabilizability again.* The fix is similar to the case of indirect MRAC: if we know the sign of b and if

$$|b| \geq b_0 \text{ (known)}$$

then we project the gradient law onto $(-\infty, -b_0]$ or $[b_0, \infty)$ (depending on the sign of b), which we know doesn't affect the properties (94).

[12] There we had a single output, while here we have full state measurements. But this only makes things easier; we can always choose a scalar output and use it to design the adaptive law. For example, $y = x_1$ gives the transfer function $\dfrac{bs + b}{s^2 + as}$. Note that there are only two parameters to estimate.

Step 2 Using (94) and well-known facts about various induced gains of exponentially stable linear systems being finite, it is possible (although not easy) to show that the perturbed system (95) is still asymptotically stable, thus $x \to 0$ as needed.

We skip Step 2. It can be found (modulo different notation and presence of extra dynamics) in [Ioannou-Sun, pp. 481–482]. (Relies on a swapping lemma.)

We see that the basic idea behind the analysis is rather simple:

- The controller stabilizes the design model for each frozen value of the parameters
- The adaptive law guarantees slow adaptation
- Theorem 10 implies stability of the time-varying closed-loop design model
- The real closed-loop system has extra terms coming from parameter estimation errors, which are small in an appropriate sense by the properties of the adaptive law
- A perturbation argument finishes the stability proof

We also learn the following lessons from this: the control design should really be based on the design model, and it should be robust with respect to the errors between the design model and the real model. This suggests a departure from the certainty equivalence principle, which doesn't address such robustness. We already saw this idea in the nonlinear context (ISS).

8.3 Detectability

The above design is not completely satisfactory. First, Step 2 is messy (we didn't give it). Second, adaptive laws based on normalization are somewhat difficult to analyze and do not extend well to nonlinear plants.

Recall that in indirect MRAC we had a different, estimator-based adaptive law with no normalization, and the estimator coincided with the reference model in closed loop. (We had a reference signal there but it's not important, we can set it to 0.) It also guaranteed slow adaptation speed in the sense of (94), and in addition we had $\dot{\hat{a}}, \dot{\hat{b}} \to 0$ and $e = y - y_m \to 0$ where y_m was the state of the estimator/model, something we don't get when we use normalization.

This suggests the following alternative strategy:

- Use an unnormalized estimator-based adaptive law
- Show stability of the estimator in closed loop using Theorem 10
- Use $e \to 0$ to prove stability of the actual plant

We don't pursue this alternative approach in detail here (mainly because the design we used for indirect MRAC doesn't generalize very well beyond plants of relative degree 1). This reasoning will be taken up later in the course (see switching adaptive control).

However, we want to plant the seed now and discuss the last step above. In the full measurement case ($y = x$), which includes the scalar case we studied in indirect MRAC, convergence of the estimator state $\hat{x}$ plus convergence of e immediately give convergence of the plant state $x = \hat{x} - e$. In general, however, e is an output of lower dimension than x and $\hat{x}$. The relevant concept then is *detectability* with respect to this output. This is a refinement of our earlier reasoning based on observability. For now we'll just make a brief general remark on detectability, and will later revisit it and will see its importance more clearly.

Suppose we write the overall closed-loop adaptive control system as

$$\begin{aligned} \dot{x} &= A_{\hat{\theta}(t)}x \\ e &= C_{\hat{\theta}(t)}x \end{aligned}$$

Here x is the combined state of plant, controller, and estimator. As usual, $\hat{\theta}$ is the vector of parameter estimates and e is the output estimation error. The above form is completely general, as long as the plant, controller, and estimator dynamics are linear.

$\longrightarrow$ The adaptive/tuning law, which is typically nonlinear, is not a part of the system dynamics.

Assume that the above system is *detectable* for each fixed value of $\hat{\theta}$. (This property is sometimes also called *tunability* [Morse].)

Recall: detectability means stability of unobservable modes, meaning that if the output equals 0 (or even just converges to 0) then the state converges to 0. Note that this is strictly weaker than observability. For now we're talking about detectability of LTI systems, i.e., detectability of each fixed pair $(A_{\hat{\theta}(t)}, C_{\hat{\theta}(t)})$.

An equivalent characterization of detectability is the existence (for each frozen value of $\hat{\theta}$) of an *output injection* matrix, $L_{\hat{\theta}}$, such that $A_{\hat{\theta}} - L_{\hat{\theta}}C_{\hat{\theta}}$ is Hurwitz.

We can take $A(\cdot)$ and $C(\cdot)$ to be C^1 in $\hat{\theta}$. Then can show that $L(\cdot)$ is also C^1 in $\hat{\theta}$.

Suppose that the adaptive law ensures that:

1) $\hat{\theta}$ is bounded

2) $\dot{\hat{\theta}}$ is in L_2, or converges to 0, or is small in some other sense so that the hypotheses of Theorem 10 are fulfilled

3) e is in L_2, or converges to 0, or is some other *zeroing signal*, i.e., signal which, when injected into an exponentially stable linear system, produces a state converging to 0

As we said, one example of an adaptive law with these properties is the one we used for indirect MRAC.

Combining this with the detectability assumption, we get stability immediately:

Rewrite the x-dynamics as

$$\dot{x} = (A_{\hat{\theta}(t)} - L_{\hat{\theta}(t)} C_{\hat{\theta}(t)})x + L_{\hat{\theta}(t)} e$$

The homogeneous part is exponentially stable by Theorem 10. (Here we are using the facts that $dA/d\hat{\theta}, dC/d\hat{\theta}, dL/d\hat{\theta}$ are continuous as functions of $\hat{\theta}$, and $\hat{\theta}$ is bounded.)

e is a zeroing signal, hence so is $L_{\hat{\theta}(t)} e$ (because L is continuous with respect to $\hat{\theta}$ and $\hat{\theta}(t)$ is bounded).

Therefore, $x \to 0$ and we are done.

This argument follows the same logic as in the previous subsection, but is much cleaner (especially the last step).

$\longrightarrow$ The above argument shows that to get stability, we should design the adaptive system so that it is detectable with respect to the output estimation error, for each frozen value of the parameter estimates.

Detectability is not only sufficient for stability, but also necessary, in the following sense. Suppose that for some value $\hat{\theta}$, the system is not detectable. This means that we can have $e \equiv 0$ but $x \not\to 0$. Of course, if $\hat{\theta}$ will change, we don't necessarily have a problem. But if $\hat{\theta}$ is an *equilibrium value* of the adaptive law, i.e., if $e \equiv 0$ will make $\hat{\theta}$ stuck at this value, then x will never converge. So, we must have detectability for all equilibrium values of the adaptive law.

9 Switching adaptive control

References: see papers on my website under the category "Supervisory Control" which build on previous work by Morse, Hespanha, and others. In particular, the paper "Overcoming the limitations of adaptive control by means of logic-based switching" (2003) contains direct comparison with traditional adaptive control. Also Chapter 6 of my book "Switching in Systems and Control" contains a detailed tutorial treatment of the subject.

Switched systems and switching control were already mentioned a couple of times earlier in the course. In particular, we looked at an example of a switched system when discussing stability of time-varying systems (Section 8.1) and we briefly discussed a switching control law when we studied universal regulators (at the end of Section 3.1.3).

Switching adaptive control aims to solve the same problems as traditional adaptive control, discussed so far in this course. It also uses many of the same concepts and ideas. The primary difference is that instead of using continuous tuning laws to define the control gains, it relies on *logic-based switching* among a family of candidate controllers.

We saw that continuous tuning/estimation has some limitations. To begin, we need a nice parameterization of the plant model. In fact, we only considered cases where the unknown parameters enter *linearly*. One reason we needed this is that for the gradient law, this guaranteed that the cost function was convex in $\hat{\theta}$. Without this property, the gradient law would not give any useful convergence results.

Another issue that we encountered was *loss of stabilizability*. Continuous tuning can take us to places in the parameter space where no stabilizing control gains exist. To overcome this, we had to use some form of projection (which requires a priori information about unknown parameters).

Switching adaptive control aims to lift these restrictions by abandoning continuous tuning altogether, and instead updating the controller gains in a discrete (*switched*) fashion. This gives us greater design flexibility.

Lecture 20 (Wed Nov 11)

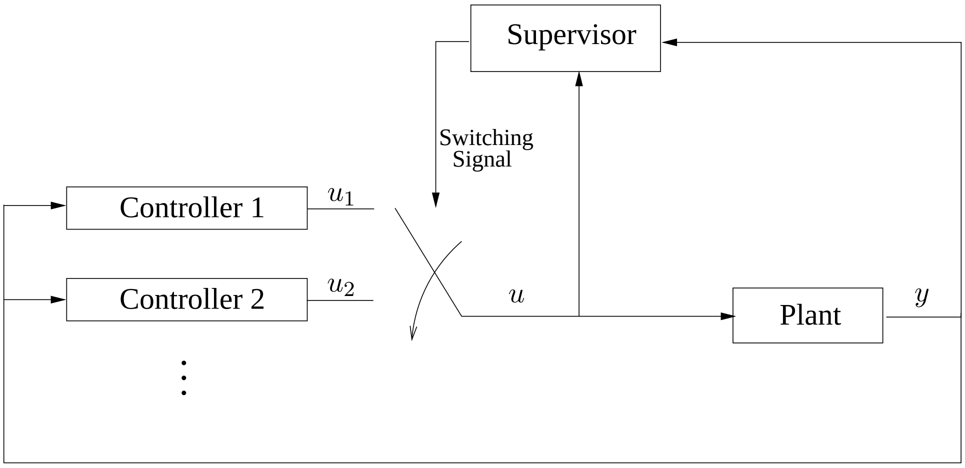

⟶ The above figure suggests that the set of controllers is discrete (even finite). This is an option that was not possible in continuous adaptive control. However, we can still have a continuum of controllers and pick controllers from this continuous family one at a time.

⟶ The unknown parameters in the plant may take values in a discrete or a continuous set. Intuitively speaking, we should have one controller for each possible value of the unknown parameter, and we'll usually assume that this is the case. However, we don't need to have exact correspondence between the two parameter sets (plant parameters and controller indices). For example, we can try to "cover" a continuous set of plant parameter values (such as a ball) with a finite number of controllers (each of which is robust enough to stabilize nearby plants).

Notation: We will write $\mathcal{P}$ for the set in which the unknown parameters take values. We assume that this is a compact set. The vector of unknown parameters itself will be denoted by p^*, and its estimates will be denoted by p. This is the same as what we earlier denoted as θ and $\hat{\theta}$, respectively, but this new notation will be more convenient (because we will frequently use p as subscripts) and it is more consistent with the switching adaptive control literature.

As in Section 7, we want to have a modular design, i.e., formulate separate design objectives on the controllers and on the supervisor.

9.1 The supervisor

We now assume for the moment that the controllers are given (we'll specify later what properties they should have) and discuss the architecture of the decision-making block, which we call the supervisor. The basic components of the supervisor are given in the following figure.

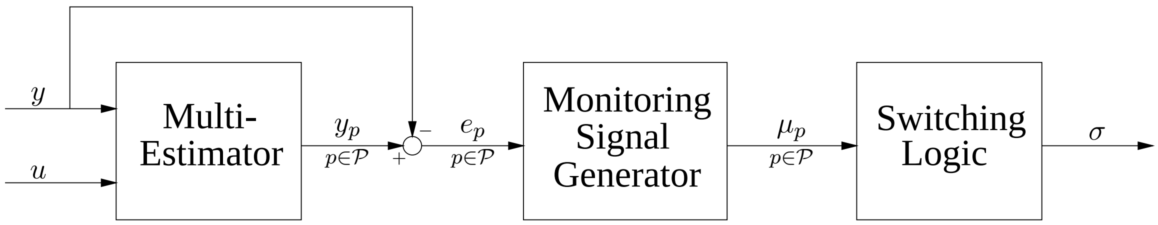

9.1.1 Multi-estimator

The scheme we discuss here is estimator-based, and is similar to the estimator equations we saw earlier. The difference is that we will no longer design a continuous tuning law to drive the output estimation error to 0, but instead we will generate a *family* of estimates

$$y_p, \ p \in \mathcal{P}$$

and the corresponding output estimation errors

$$e_p := y_p - y, \qquad p \in \mathcal{P}$$

and will *pick the smallest one* at each time instant (roughly speaking).

$\longrightarrow$ The basic idea is to design the multi-estimator so that e_{p^*}, i.e., the error corresponding to the true parameter value, is small. Usually we cannot guarantee anything about the other, "wrong" estimation errors (there is no a priori reason for them to be small). Thus, the smallness of e_p indicates the likelihood that $p = p^*$. In other words, it seems intuitively reasonable (although not yet justified formally in any way) to pick as a current estimate of p^* the index of the smallest estimation error.

This property is basically the same as what we had for the estimators we had earlier: in case when the parameter estimates match the true values, e is small—e.g., converges to 0 exponentially at a prescribed rate. (See the estimator for Example 6 in Section 6.2, which we saw again in indirect MRAC.)

Example 10 Consider the scalar example

$$\dot{y} = y^2 + p^* u$$

where p^* is an unknown element of some set $\mathcal{P} \subset \mathbb{R}$ containing both positive and negative values. We can let the estimator equations be

$$\dot{y}_p = -a_m(y_p - y) + y^2 + pu, \qquad p \in \mathcal{P} \tag{96}$$

where $a_m > 0$. Then the estimation error $e_{p^*} = y_{p^*} - y$ satisfies

$$\dot{e}_{p^*} = -a_m e_{p^*}$$

and hence converges to 0 exponentially fast, for an arbitrary control u.

(It is useful to compare this with the estimator

$$\dot{\hat{y}} = -a_m(\hat{y} - y) + y^2 + \hat{p}u$$

which we would have used in conventional adaptive control.)

Note that if we add a disturbance to the plant:

$$\dot{y} = y^2 + p^* u - d$$

then we have

$$\dot{e}_{p^*} = -a_m e_{p^*} + d$$

hence e_{p^*} exponentially converges to d/a_m, not to 0. In this case, the steady-state value of the error e_{p^*} is determined by the size of the disturbance.

One concern is that realizing the multi-estimator simply as a parallel connection of individual estimator equations for $p \in \mathcal{P}$ is not efficient and actually impossible if $\mathcal{P}$ is an infinite set. The estimator equations (96) can be implemented differently as follows. Consider the system

$$\begin{aligned} \dot{z}_1 &= -a_m z_1 + a_m y + y^2 \\ \dot{z}_2 &= -a_m z_2 + u \end{aligned} \tag{97}$$

together with the outputs

$$y_p := z_1 + p z_2, \qquad p \in \mathcal{P} \tag{98}$$

The two-dimensional system (97) produces the same signals as does the (possibly infinite-dimensional) system (96):

$$\dot{y}_p = \dot{z}_1 + p\dot{z}_2 = -a_m z_1 + a_m y + y^2 - p a_m z_2 + pu \underset{\text{by (98)}}{=} -a_m y_p + a_m y + y^2 + pu$$

This idea is known as *state sharing*. The family of signals (98) is of course still infinite, but at each particular time we can look any one of them up or perform mathematical operations—such as computing the minimum—with the entire family. □

9.1.2 Monitoring signal generator

Rather than basing decisions on the instantaneous values of the output estimation errors, we would like to take their past behavior into account. Thus we need to implement an appropriate filter, which we call the monitoring signal generator. This is a dynamical system whose inputs are the estimation errors and whose outputs

$$\mu_p, \ p \in \mathcal{P}$$

are suitably defined integral norms of the estimation errors, called *monitoring signals*. For example, we can simply work with the squared L_2 norm:

$$\mu_p(t) := \int_0^t |e_p(s)|^2 ds, \qquad p \in \mathcal{P} \tag{99}$$

These monitoring signals can be generated by the differential equations

$$\dot{\mu}_p = |e_p|^2, \qquad \mu_p(0) = 0, \quad p \in \mathcal{P} \tag{100}$$

If we don't want the signals μ_p to grow unbounded, we can introduce a "forgetting factor":

$$\mu_p(t) := \int_0^t e^{-\lambda(t-s)} |e_p(s)|^2 ds$$

with $\lambda > 0$, which can be implemented via

$$\dot{\mu}_p = -\lambda \mu_p + |e_p|^2, \qquad \mu_p(0) = 0, \quad p \in \mathcal{P} \tag{101}$$

Again, we do not want to generate each monitoring signal individually. The idea of state sharing can be applied here as well. To see how this works, let us revisit the multi-estimator for Example 10, given by (97) and (98). Each estimation error can be equivalently expressed as

$$e_p = z_1 + pz_2 - y$$

so that we have

$$e_p^2 = (z_1 - y)^2 + 2pz_2(z_1 - y) + p^2 z_2^2, \qquad p \in \mathcal{P}$$

If we now define the monitoring signal generator via

$$\begin{aligned} \dot{\eta}_1 &= (z_1 - y)^2 \\ \dot{\eta}_2 &= 2z_2(z_1 - y) \\ \dot{\eta}_3 &= z_2^2 \\ \mu_p &= \eta_1 + p\eta_2 + p^2\eta_3, \qquad p \in \mathcal{P} \end{aligned} \tag{102}$$

then the equations (100) still hold. By now you should be able to see why this works and how it can be extended to (101).

9.1.3 Switching logic

This is a dynamical system whose inputs are the monitoring signals μ_p, $p \in \mathcal{P}$ and whose output is a piecewise constant *switching signal* σ. The switching signal determines the actual control law

$$u = u_\sigma$$

applied to the plant, where u_p, $p \in \mathcal{P}$ are the control signals generated by the candidate controllers.

$\longrightarrow$ We don't actually need to physically generate the out-of-the-loop control signals $u_p, p \neq \sigma(t)$. The diagram with a bank of controllers is just for illustration. What we're really implementing is a *switching controller*.

$\longrightarrow$ The controllers can be parameterized by a different set $\mathcal{Q}$, which can for example be a subset of $\mathcal{P}$. This can be done by defining a *controller assignment map* from $\mathcal{P}$ to $\mathcal{Q}$. To keep things simple, we will assume that $\mathcal{Q} = \mathcal{P}$. But we'll usually write q and not p for controller indices, to distinguish them from the plant parameters and to remind us of the more general situation. We'll also sometimes denote the controllers themselves as $\mathbb{C}_q$, $q \in \mathcal{P}$.

Basic idea of the switching logic:

$$\sigma(t) := \arg\min_{p\in\mathcal{P}} \mu_p(t)$$

This is essentially what we'll do, but we'll only update the value of σ from time to time, rather than continuously. This discrete update strategy for σ can be either time-based (use a fixed time interval between updates—*dwell time*) or event-based (update when the difference between the old minimum and the new one gets large enough—*hysteresis*). More details on this later.

Justification–?

$$\mu_q \text{ small} \quad\Longrightarrow\quad e_q \text{ small} \quad \underset{\text{WRONG}}{\Longrightarrow} \quad q = p^* \quad\Longrightarrow\quad q\text{-th controller stabilizes the plant}$$

This is tempting, but the second implication is not known to hold. In fact, what we have is its *converse*.

We saw this issue before in continuous adaptive control: the estimation scheme drives e to 0, and we plug the estimate $\hat\theta$ into the controller as if it were θ (*certainty equivalence*). However, even if $e \to 0$, we don't know that $\hat\theta \to \theta$, so a separate justification is required why this works. Before, we were using Lyapunov functions to prove stability. Here the reasoning will be different, based on *detectability*.

In fact, let's drop the wrong implication and state directly the property we want:

$$e_q \text{ small} \quad \underset{\text{want this}}{\Longrightarrow} \quad q\text{-th controller stabilizes the plant}$$

This is precisely the detectability property, which we already discussed in Section 8.3. And it must be true *for all* q since we don't know the behavior of σ a priori.

More precisely, this is detectability of the supervisory control system with the q-th controller in the loop, with respect to the q-th output estimation error. Assume that this system is linear, and write it as

$$\begin{aligned}\dot x &= A_q x\\ e_q &= C_q x\end{aligned}$$

Recall (we'll need this later) that detectability is equivalent to the existence (for each frozen value of q) of an *output injection* matrix, L_q, such that $A_q - L_q C_q$ is Hurwitz. Use this to rewrite the system as

$$\dot{x} = (A_q - L_q C_q)x + L_q e_q$$

It is then clear that small e_q (bounded, convergent to 0) implies the same property for x.

9.2 Example: linear systems

Consider a linear plant

$$\begin{aligned} \dot{x} &= A_{p^*}x + B_{p^*}u \\ y &= C_{p^*}x \end{aligned} \tag{103}$$

where $x \in \mathbb{R}^n$, $u \in \mathbb{R}^m$, $y \in \mathbb{R}^k$, $\{A_p, B_p, C_p : p \in \mathcal{P}\}$ is a given finite family of matrices, and $p^* \in \mathcal{P}$ is unknown. We assume that each system in this family is controllable and observable (or at least stabilizable and detectable).

Problem: regulate x to 0 using output feedback.

We haven't discussed estimator design for the case of output measurements (we've only considered full state measurements or scalar systems before). However, we can use the standard Luenberger observer to design the multi-estimator:

$$\begin{aligned} \dot{x}_p &= (A_p - L_p C_p)x_p + B_p u + L_p y \\ y_p &= C_p x_p \end{aligned} \qquad p \in \mathcal{P}$$

where each L_p is an output injection matrix such that $A_p - L_p C_p$ is Hurwitz. It follows that the estimation error $e_{p^*} = y_{p^*} - y$ converges to zero exponentially fast, regardless of the control u that is applied (just subtract the two right-hand sides when $p = p^*$).

We can also use the above observers to design the candidate control laws:

$$u_p = -K_p x_p, \qquad p \in \mathcal{P} \tag{104}$$

where the matrices K_p are such that $A_p - B_p K_p$ are Hurwitz for each $p \in \mathcal{P}$.

Let the monitoring signals be the L_2 norms of the output estimation errors as in (99), i.e., generate them by the differential equations (100).

Lecture 21 (Mon Nov 16)

It remains to define the switching signal $\sigma : [0, \infty) \to \mathcal{P}$, which will give us the switching controller $u(t) = -K_{\sigma(t)} x_{\sigma(t)}$. One way to do this is by means of the so-called *hysteresis switching logic*, see figure.

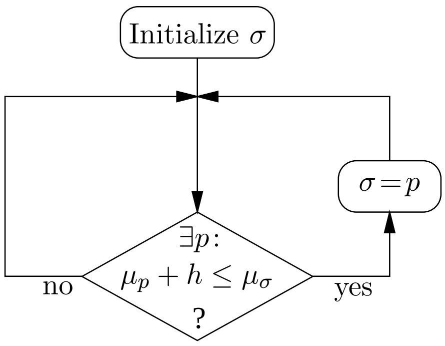

This switching logic works as follows. Fix a positive number h called the *hysteresis constant.* Set $\sigma(0) = \arg\min_{p \in \mathcal{P}} \mu_p(0)$. Now, suppose that at a certain time σ has just switched to some $q \in \mathcal{P}$. The value of σ is then held fixed until we have $\min_{p \in \mathcal{P}} \mu_p(t) + h \leq \mu_q(t)$. If and when that happens, we set σ equal to $\arg\min_{p \in \mathcal{P}} \mu_p(t)$. When the indicated arg min is not unique, break the tie arbitrarily.

$\longrightarrow$ In Example 10, for each fixed time t, μ_p is a quadratic polynomial in p given by the formula (102), hence the value of p that minimizes $\mu_p(t)$ is either a real root of $\frac{\partial \mu_p}{\partial p}(t)$ or a boundary point of $\mathcal{P}$. Thus for that example, the on-line minimization procedure required for implementing the switching logic is relatively straightforward.

Lemma 11 *The switching stops in finite time, i.e., there exists a time T^* and an index $q^* \in \mathcal{P}$ such that $\sigma(t) = q^* \in \mathcal{P}$ for all $t \geq T^*$. Moreover, $e_{q^*} \in L_2$.*

Proof. Since e_{p^*} converges to 0 exponentially, the formula (99) implies that $\mu_{p^*}(t)$ is bounded from above by some number K for all $t \geq 0$. In addition, all monitoring signals μ_p are nondecreasing by construction. Using these two facts and the definition of the hysteresis switching logic, we now prove that the switching must stop in finite time. Indeed, each μ_p has a limit (possibly ∞) as $t \to \infty$. Since $\mathcal{P}$ is finite, there exists a time T such that for each $p \in \mathcal{P}$ we either have $\mu_p(T) > K$ or $\mu_p(t_2) - \mu_p(t_1) < h$ for all $t_2 > t_1 \geq T$. Then for $t \geq T$ at most one more switch can occur. We conclude that there exists a time T^* such that $\sigma(t) = q^* \in \mathcal{P}$ for all $t \geq T^*$, and μ_{q^*} is bounded. Finally, $e_{q^*} \in L_2$ by (99). $\square$

After the switching stops, the closed-loop system (excluding out-of-the-loop signals) can be written as

$$\begin{aligned}\begin{pmatrix}\dot{x}\\ \dot{x}_{q^*}\end{pmatrix} &= \overline{A}\begin{pmatrix}x\\ x_{q^*}\end{pmatrix}\\ e_{q^*} &= \overline{C}\begin{pmatrix}x\\ x_{q^*}\end{pmatrix}\end{aligned} \tag{105}$$

where

$$\overline{A} := \begin{pmatrix}A_{p^*} & -B_{p^*}K_{q^*}\\ L_{q^*}C_{p^*} & A_{q^*} - L_{q^*}C_{q^*} - B_{q^*}K_{q^*}\end{pmatrix}$$

and

$$\overline{C} := \begin{pmatrix}-C_{p^*} & C_{q^*}\end{pmatrix}.$$

If we let

$$\overline{L} := \begin{pmatrix}-L_{p^*}\\ -L_{q^*}\end{pmatrix}$$

then it is straightforward to check that

$$\overline{A} - \overline{L}\,\overline{C} = \begin{pmatrix}A_{p^*} - L_{p^*}C_{p^*} & -B_{p^*}K_{q^*} + L_{p^*}C_{q^*}\\ 0 & A_{q^*} - B_{q^*}K_{q^*}\end{pmatrix}.$$

The matrix on the right-hand side is Hurwitz, which shows that the system (105) is detectable with respect to e_{q^*}.

It now remains to apply the standard output injection argument. Namely, write

$$\begin{pmatrix}\dot{x}\\ \dot{x}_{q^*}\end{pmatrix} = (\overline{A} - \overline{L}\,\overline{C})\begin{pmatrix}x\\ x_{q^*}\end{pmatrix} + \overline{L}e_{q^*}$$

and observe that x and x_{q^*} converge to zero in view of stability of $\overline{A} - \overline{K}\,\overline{C}$ and the fact that $e_{q^*} \in L_2$.

This reasoning is similar to what we had for slowly time-varying system, but there's no time variation here after the switching has stopped so this is easier.

The above problem is rather special, and the solution and the method of proof have several drawbacks:

- The Luenberger-based multi-estimator that we gave only works when $\mathcal{P}$ is a finite set. It is not suitable for state-shared implementation. However, there are standard results in linear identification that can be used to design state-shared multi-estimators. (We will not discuss this.)

- Our controller design and multi-estimator design were coupled because both relied on the Luenberger observer. However, the particular choice of candidate control laws (104) is just one example, and can be easily changed. Assume, for example, that every system in the family to which the plant (103) belongs is stabilizable by a static linear output feedback. In other words, assume that for each $p \in \mathcal{P}$ there exists a matrix K_p such that the eigenvalues of $A_p - B_pK_pC_p$ have negative real parts. A straightforward modification of the above argument shows that if we keep the estimators as they are but replace the control laws (104) by $u_p = -K_p y$, we still achieve state regulation.

- Detectability of the feedback connection of the plant with the q^*-th controller seems to come out of the blue. Actually, it is a natural consequence of the fact that this controller stabilizes the corresponding plant model (for $p^* = q^*$). Here is another proof of detectability which makes this clearer. Consider first the dynamics of x_{q^*} after switching stops, i.e., under the action of the q^*-th controller:

$$\dot{x}_{q^*} = (A_{q^*} - L_{q^*}C_{q^*})x_{q^*} - B_{q^*}K_{q^*}x_{q^*} + L_{q^*} \underbrace{C_{p^*}x}_{C_{q^*}x_{q^*} - e_{q^*}} = (A_{q^*} - B_{q^*}K_{q^*})x_{q^*} - L_{q^*}e_{q^*}$$

  Suppose that $e_{q^*} \to 0$ (this is just for showing detectability, it might not actually be true). Then $x_{q^*} \to 0$ because $A_{q^*} - B_{q^*}K_{q^*}$ is Hurwitz. This implies that $y = C_{q^*}x_{q^*} - e_{q^*} \to 0$, and also $u = -K_{q^*}x_{q^*} \to 0$. Since the plant is assumed to be detectable and $u, y \to 0$, we have $x \to 0$ (apply an output injection argument again, this time to the plant only). We have just shown that the (x, x_{q^*})-system, driven by the q^*-th controller, is detectable with respect to e_{q^*} as desired. Thus we know that the output injection matrix $\overline{L}$, which we so magically found earlier, must exist.

In any case, the above example was a useful illustration of the main ideas, and we will now take another pass through the approach to refine and generalize it.

9.3 Modular design objectives and analysis steps

We see that detectability is a key concept here.

Up to now, we discussed detectability of the closed-loop system for each *fixed* controller with respect to the corresponding estimation error. In the above example, we applied this to the q^*-th controller, where q^* is the index at which the switching stops.

However, if the switching doesn't stop, we need *switched detectability*. In the linear case, this means that if we write the switched closed-loop system as

$$\begin{aligned} \dot{x} &= A_\sigma x \\ e_\sigma &= C_\sigma x \end{aligned} \tag{106}$$

then we want it to be detectable with respect to e_σ. In other words, e_σ being small ($\to 0$) should imply that x is small ($\to 0$). Here x is the state of the plant, the multi-estimator, and the active controller (at each time).

Does this switched detectability property follow from detectability for each frozen value of σ (which we assumed earlier)?

No. This is similar to the stability issue, which for time-varying or switched systems doesn't follow from stability of individual fixed subsystems among which we are switching.

In fact, the connection with stability is quite direct. Recall that detectability for each fixed index q is equivalent to the existence of a matrix L_q such that $A_q - L_q C_q$ is Hurwitz. Use this to rewrite the switched system as

$$\dot{x} = (A_\sigma - L_\sigma C_\sigma)x + L_\sigma e_\sigma \tag{107}$$

We know that the switched system

$$\dot{x} = (A_\sigma - L_\sigma C_\sigma)x$$

may not be stable even if each $A_q - L_q C_q$ is Hurwitz. Some further properties of σ are required.

We saw this in stability of time-varying systems. One option is to require that σ be *slow enough.* It no longer makes sense to bound its derivative (as in Chapter 8) because σ is not differentiable. But there are several possible slow-switching conditions that work:

- Switching stops in finite time (as in the above example). Then stability is obvious. But this condition is quite strong.
- There exists a sufficiently large *dwell time* τ_D, i.e., the time between any two consecutive switches is lower-bounded by τ_D.
- There exists a sufficiently large *average dwell time* τ_{AD}. This means that the number of switches on any interval $(t, T]$, which we denote by $N_\sigma(T, t)$, satisfies

  $$N_\sigma(T, t) \le N_0 + \frac{T - t}{\tau_{AD}}$$

  Here $N_0 \ge 1$ is an arbitrary number (it cannot depend on the choice of the time interval).

  For example, if $N_0 = 1$, then σ cannot switch twice on any interval of length smaller than τ_{AD}. This is exactly the dwell time property. Note also that $N_0 = 0$ corresponds to the case of no switching, since σ cannot switch at all on any interval of length smaller than τ_{AD}. In general, if we discard the N_0 "extra" switches, then the average time between consecutive switches is at least τ_{AD}. Average dwell time is more general than dwell time, because it allows us to switch fast when necessary and then compensate for it by switching sufficiently slowly later.

There is a large literature on stability of switched systems. It is known that any of the above conditions guarantees that stability is preserved under switching. The last result (on average dwell time) is due to Hespanha.

Now, suppose σ does satisfy one of the above conditions. Then we can conclude detectability of (106), i.e., we know that $x \to 0$ if $e_\sigma \to 0$, in view of (107).

But for this to be useful, we need to know that e_σ is small. We hope that the switching logic will somehow guarantee this, even if the switching doesn't stop.

We are now ready to state, at least qualitatively, four main design objectives placed on the individual components of the supervisory control system:

Matching At least one of the estimation errors is small

Detectability For each fixed controller, the closed-loop system is detectable through the corresponding estimation error

Bounded Error Gain The signal e_σ is bounded in terms of the smallest of the estimation errors

Non-Destabilization The switched closed-loop system is detectable through e_σ provided that detectability holds for every frozen value of σ

The Matching property is a requirement imposed on the multi-estimator design. The estimation error that we can hope to be small is e_{p^*}, where p^* is the true value of the unknown parameter. The above somewhat more vague statement is a bit more general.

In Example 10, we could get $e_{p^*} \to 0$ when there is no disturbance, and we could get e_{p^*} bounded when there is a bounded disturbance (regardless of the control law being applied).

For linear systems, there exists a theory for designing multi-estimators with such properties. For some classes of nonlinear systems, this can be done along the lines of what we saw in Example 10 and earlier: put terms on the right-hand side of the multi-estimator that *match* the right-hand side of the plant, and add linear damping.

We will not discuss the Matching property further.

The Detectability property is a requirement imposed on the candidate controllers. It is interesting and we'll discuss it further below.

The last two properties are requirements placed on the switching logic. We'll also discuss them a little bit below.

It is not difficult to see now, at least conceptually, how the above properties of the various blocks of the supervisory control system can be put together to analyze its behavior.

Analysis:

$$\left.\begin{array}{l}\text{Matching} + \text{Bounded Error Gain} \implies e_\sigma \text{ is small} \\ \text{Detectability} + \text{Non-Destabilization} \implies \text{detectability with respect to } e_\sigma\end{array}\right\} \implies x \text{ is small}$$

Lecture 22 (Wed Nov 18)

9.3.1 Achieving detectability

The first figure shows the closed-loop system which we want to be detectable, with e_q viewed as output.

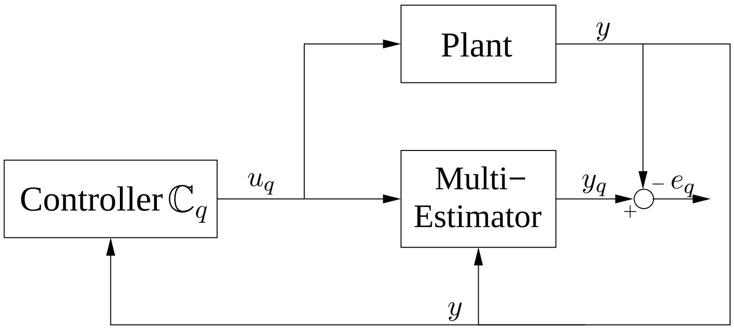

The second figure shows an equivalent but more convenient representation of the same system.

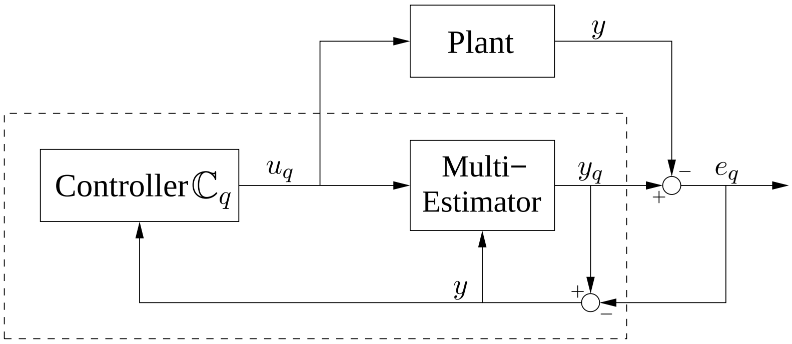

The system inside the box is called the *injected system.* It is the connection of the q-th controller with the multi-estimator, and e_q is an input injected into it.

$\longrightarrow$ The reason this transformation is useful is because, as we already discussed earlier, the control design should really be based on the design model (multi-estimator in the present setting) and not on the unknown plant itself.

We now state, in very informal terms, a representative "theorem" which gives sufficient conditions for detectability.

Theorem 12 *Let $x_{\mathbb{P}}, x_{\mathbb{C}}, x_{\mathbb{E}}$ be the states of the plant, the controller, and the multi-estimator, respectively. Assume that:*

1) The injected system satisfies

$$e_q \textit{ small} \implies x_{\mathbb{C}}, x_{\mathbb{E}} \textit{ small}$$

2) The plant satisfies

$$u_q, y \textit{ small} \implies x_{\mathbb{P}} \textit{ small}$$

Then the overall closed-loop system is detectable, i.e.,

$$e_q \textit{ small} \implies x_{\mathbb{P}}, x_{\mathbb{C}}, x_{\mathbb{E}} \textit{ small}$$

"Proof"—almost immediate:

If e_q is small, then 1) guarantees that $x_{\mathbb{C}}$ and $x_{\mathbb{E}}$ are small.

u_q and y_q are then small since they are functions of $x_{\mathbb{C}}$ and $x_{\mathbb{E}}$.

Hence $y = y_q - e_q$ is also small.

Finally, 2) guarantees that $x_{\mathbb{P}}$ is small. □

The above result can be made completely rigorous. Property 1) of the injected system can be formalized using the ISS notion, i.e., by saying that the controller input-to-state stabilizes the multi-estimator. Thus the material of Section 7.2 is directly relevant here.

⟶ Actually, here things are even simpler because the estimation error e_q, with respect to which we want to achieve ISS, is known and can be used by the controller (while in Section 7.2 the disturbance d was assumed to be unknown).

⟶ If the dynamics are linear, then as we know this ISS property is automatic from the internal asymptotic stabilization.

Property 2) is detectability of the plant, and can be formalized in a way similar to ISS (input-output-to-state stability). A formal proof then requires some manipulations with class $\mathcal{K}_\infty$ and $\mathcal{KL}$ functions.

Actually, ISS is an overkill. Can use a weaker property, called integral-ISS, instead. This is suitable because estimation errors are typically small in an integral sense, since we integrate them to construct monitoring signals. The injected system we get for Example 10 with the obvious choice of the control laws

$$u_p = -\frac{y^2 + y}{p}, \qquad p \in \mathcal{P}$$

is not ISS, but it can be shown to be integral ISS.

Here is another representative result:

Theorem 13 *Assume that:*

1) The injected system satisfies

$$e_q \textit{ small} \implies y_q \textit{ small}$$

2) The plant satisfies

$$y \text{ small} \Longrightarrow x_{\mathbb{P}}, u_q \text{ small}$$

3) The controller satisfies

$$y, u_q \text{ small} \Longrightarrow x_{\mathbb{C}} \text{ small}$$

4) The multi-estimator satisfies

$$y, u_q, y_q \text{ small} \Longrightarrow x_{\mathbb{E}} \text{ small}$$

Then the overall closed-loop system is detectable, i.e.,

$$e_q \text{ small} \Longrightarrow x_{\mathbb{P}}, x_{\mathbb{C}}, x_{\mathbb{E}} \text{ small}$$

1) is input-to-*output* stability of the injected system. This is weaker than 1) in the previous theorem.

2) is essentially[13] the *minimum-phase* property of the plant. This is stronger than 2) in the previous theorem. The requirement that the plant be minimum-phase is actually quite common in adaptive control; see [Ioannou-Sun, pp. 25, 332, 412, and elsewhere].

3), 4) are detectability of the controller and multi-estimator. These are usually reasonable assumptions.

Can "prove" this result by a simple signal-chasing argument along the lines of the previous one. Try it!

9.3.2 Achieving bounded error gain and non-destabilization

We now turn to the Bounded Error Gain and Non-Destabilization properties, which need to be enforced by the switching logic.

The Bounded Error Gain property is about making e_σ small. Its counterpart in continuous adaptive control was the smallness of the output estimation error e (in the sense of L_2 or convergence to 0). Here, the bound for e_σ is usually stated in terms of suitable integral expressions for the estimation errors, which are related to the monitoring signals.

The Non-Destabilization property, as we discussed, is about σ switching slowly enough (in the sense of dwell time, average dwell time, or termination of switching if we're lucky). Its counterpart in continuous adaptive control was slow adaptation speed ($\dot{\hat{\theta}}$ being in L_2 or converging to 0).

[13]The definition of minimum phase is stated in terms of $y \equiv 0$, while here we say "y is small." For linear plants this makes no difference. For nonlinear plants, a suitably strengthened notion of minimum phase is needed. This stronger notion relates to the usual notion of minimum-phase nonlinear system described in Section 3.2 in much the same way as ISS relates to 0-GAS.

It's not hard to see that these two properties are actually *conflicting*. To enforce bounded error gain, we want to switch to $\arg\min_p \mu_p(t)$ as quickly as possible. But for non-destabilization, we want to switch slowly enough (or even stop). So, need to find a compromise.

One option is to use the *dwell-time switching logic*, which enforces dwell time by construction.

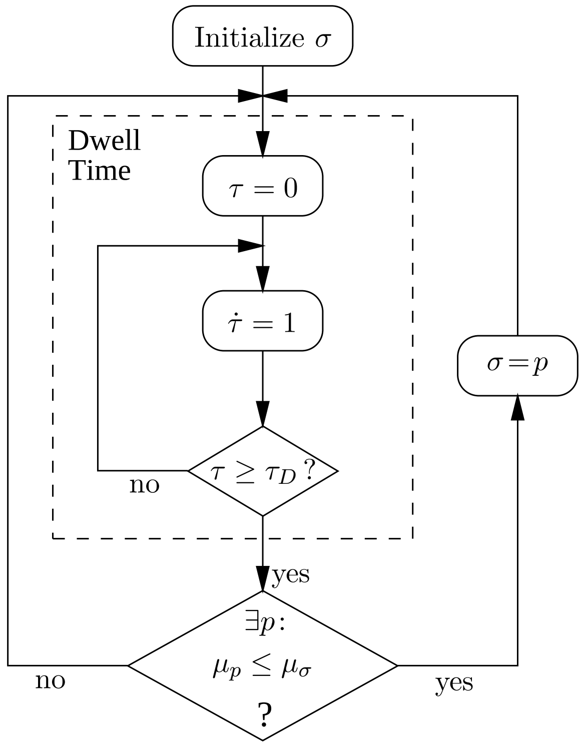

This switching logic ensures non-destabilization, and is easy to implement. Bounded error gain is harder to get, because if the switching is only allowed at event times separated by τ_D, then μ_σ is not always guaranteed to be small compared to the other monitoring signals. Although there are ways to handle this problem, we see that dwell-time switching has significant shortcomings. With a prespecified dwell time, the performance of the currently active controller might deteriorate to an unacceptable level before the next switch is permitted. If the system is nonlinear, the trajectories may even escape to infinity in finite time.

Another option (more suitable for nonlinear plants) is to use the *hysteresis switching logic*. Hysteresis means that we do not switch every time $\min_{p\in\mathcal{P}} \mu_p(t)$ becomes smaller than $\mu_\sigma(t)$, but switch only when it becomes "significantly" smaller. The threshold of tolerance is determined by a *hysteresis constant* $h > 0$.

We already used this idea in Section 9.2. The only difference is that here we are using *multiplicative* and not additive hysteresis. (This is better for technical reasons; in particular, σ doesn't change if we scale all monitoring signals by some, possibly time-varying, factor.)

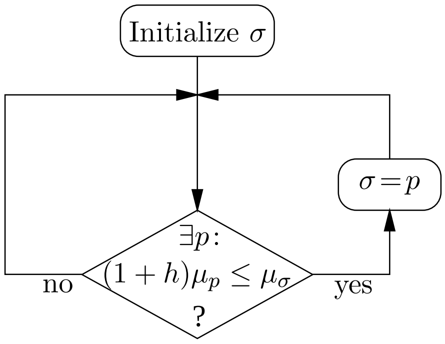

We already saw that when e_{p^*} is exponentially converging to 0 (the case of no disturbance), switching stops in finite time.

It turns out that when e_{p^*} is just bounded (the case of bounded disturbance), the switching doesn't stop but we do get *average dwell time*, which can be made large by increasing the hysteresis constant h. This is quite a remarkable result[14] due to Hespanha and Morse.

This gives the Non-Destabilization property.

The Bounded Error Gain property can also be obtained, in the form

$$\int_0^t |e_\sigma(\tau)|^2 d\tau \le |\mathcal{P}|(1+h) \min_{p\in\mathcal{P}} \int_0^t |e_p(\tau)|^2 d\tau$$

where $|\mathcal{P}|$ is the cardinality of $\mathcal{P}$.

$\longrightarrow$ This is only valid when the index set $\mathcal{P}$ is finite. When $\mathcal{P}$ is infinite, we need to consider its finite partition, and the switching logic needs to be modified.

Summary of switching adaptive control Putting together the above design ingredients:

- Multi-estimator design giving the Matching property
- Input-to-state stabilizing controller design (via Theorem 12)
- Hysteresis switching logic (giving average dwell time and bounded error gain)

one obtains quite general results for linear systems as well as some classes of nonlinear systems. The analysis proceeds along the lines outlined on page 120. It is of course not quite as simple, but arguably simpler than analysis of traditional, continuously-tuned adaptive control algorithms. The observed performance is also typically quite good. See the references given at the beginning of this chapter.

[14] In fact, the concept of average dwell time originally arose out of the study of hysteresis-based switching logics for adaptive control.

Note that the first two bullets above are reiterations of issues we've already seen in continuous adaptive control. First, we need to be able to design an estimation scheme—this places constraints on what types of systems we can deal with (and in particular, on the way in which uncertain parameters can enter the dynamics). In continuous adaptive control we typically required the parameters to enter linearly, but this is not necessary here as long as state-sharing is possible.

Second, the controllers need to provide robustness to parameter estimation errors, beyond the usual certainty equivalence stabilization assumption. We saw this type of property earlier already.

The last bullet is specific to switching adaptive control. Switching in fact allows us to overcome some difficulties associated with continuous tuning, as already discussed earlier.

Optional exercise: simulate Example 10. (The ISS condition is not easily enforceable, but the controller works anyway.)

10 Singular perturbations

Reference: [Khalil, Chapter 11]

10.1 Unmodeled dynamics

In this course, the modeling uncertainty in the plant has primarily been described by unknown parameters (*parametric uncertainty*). But this is not the only form of modeling uncertainty.

A more general formulation should allow for the presence of *dynamic* uncertainties, i.e., dynamics which are neglected at the control design stage (because they are hard to model and/or are insignificant in some sense). Constant unknown parameters are then a special case ($\dot{\theta} = 0$).

Example 11 Suppose that the plant is given by the transfer function

$$g(s) = \frac{b}{(s-a)(\varepsilon s + 1)}$$

where a, b, ε are unknown parameters, $b > 0$, and $\varepsilon > 0$ is small.

We can write this as a cascade of two systems:

$$y = \frac{b}{s-a} z, \qquad z = \frac{1}{\varepsilon s + 1} u$$

This suggests the following state-space representation:

$$\begin{aligned} \dot{y} &= ay + bz \\ \varepsilon \dot{z} &= -z + u \end{aligned} \tag{108}$$

Here we think of ε as a small "parasitic" time constant, corresponding to a very fast extra pole at $-1/\varepsilon$. It gives fast unmodeled dynamics. In the limit as $\varepsilon \to 0$ the z-dynamics disappear, and for $\varepsilon = 0$ the system reduces to

$$y = \frac{b}{s-a}u$$

or

$$\dot{y} = ay + bu \tag{109}$$

while the differential equation for z degenerates into the hard constraint

$$z = u$$

One says that the system (108) is a *singular perturbation* of (109). The perturbation is "singular" because for $\varepsilon > 0$ the dimension jumps from 1 to 2, which is very different from perturbations which simply contribute extra terms on the right-hand side of the nominal system. □

There are other tools for handling unmodeled dynamics, based on robust control theory (and small-gain theorems). We will not discuss those in this course (see [Ioannou-Sun, Sect. 8.2]).

10.2 Singular perturbations

The general singular perturbation model is

$$\begin{aligned} \dot{x} &= f(t, x, z, \varepsilon) \\ \varepsilon\dot{z} &= g(t, x, z, \varepsilon) \end{aligned}$$

where $x \in \mathbb{R}^n$ and $z \in \mathbb{R}^m$. We think of the x-system as the *slow system* and of the z-system as the *fast system*, since for ε close to 0, $\dot{z}$ becomes very large (imagine dividing the second ODE by ε).

For $\varepsilon = 0$, the second ODE becomes the algebraic constraint

$$g(t, x, z, 0) = 0$$

Suppose that we can solve this equation for z:

$$z = h(t, x)$$

Plugging this into the x-system, and keeping $\varepsilon = 0$, we obtain the *reduced system*

$$\dot{x} = f(t, x, h(t, x), 0) \tag{110}$$

Denote the solution of this system (starting from a given initial condition) as $\bar{x}(t)$.

We would like to know to what extent $\bar{x}(t)$ serves as a good approximation to the actual solution $x(t)$, i.e., to the x-component of the solution of our overall singularly perturbed system, when ε is very small.

Even if ε is small, initially $z(t)$ may be quite far from $h(t,x)$, and the reduced system will not be a valid approximation. However, since the dynamics of z are fast, we can expect that $z(t)$ will converge to $h(t,x)$ very fast, and after some (small) time t^*, z will be close to $h(t,x)$. Then, if we have initialized $\bar{x}(t)$ correctly, this initial fast transient of z will not have significant effect and $\bar{x}(t)$ will be close to $x(t)$ for all time.

But, to make sure that z becomes close to its equilibrium value $h(t,x)$, we need to have some stability property for the z-system.

Lecture 23 (Mon Nov 30)

To analyze the z-system, it is convenient to make the following two transformations. First, introduce

$$\bar{z} := z - h(t,x)$$

which has the effect of shifting the equilibrium to the origin. We have

$$\varepsilon\dot{\bar{z}} = \varepsilon\dot{z} - \varepsilon\frac{d}{dt}h(t,x) = g(t,x,\bar{z}+h(t,x),\varepsilon) - \varepsilon\frac{d}{dt}h(t,x)$$

Second, introduce the "stretched time"

$$\tau := \frac{t}{\varepsilon}$$

We get

$$\frac{d\bar{z}}{d\tau} = \varepsilon\frac{d\bar{z}}{dt} = g(t,x,\bar{z}+h(t,x),\varepsilon) - \varepsilon\frac{d}{dt}h(t,x) \tag{111}$$

Setting $\varepsilon = 0$ in the above equation, we obtain the auxiliary system

$$\frac{d\bar{z}}{d\tau} = g(t,x,\bar{z}+h(t,x),0) \tag{112}$$

It is this system that we want to be stable. It has an equlibrium at $\bar{z} = 0$ by construction. In this system, t and x are actually *fixed parameters*. Indeed, we have

$$t = \varepsilon\tau$$

and so by setting $\varepsilon = 0$ we are freezing the original time t (and hence $x(t)$ also).

$\longrightarrow$ Here t and τ give two completely different time scales.

Assumption 1: The system (112) is exponentially stable, uniformly in (t,x). This means that we have an exponential estimate

$$|\bar{z}(\tau)| \le ce^{-\lambda\tau}|\bar{z}(0)|$$

which holds for each fixed value of t and x in (112).

When ε is small but positive, t and x will vary—but slowly. In other words, in the τ time scale, we'll have a slowly time-varying system. And the ε-dependent terms will enter it. Using results on stability of slowly-time varying systems (cf. Chapter 8) one can show that the actual $\bar{z}$, described by (111) for $\varepsilon > 0$, will eventually become of order ε, and $\bar{x}$ will indeed be a good approximation of x (again up to terms of order ε).

More precisely, the statement is as follows. Pick some times $T > t^* > 0$. Let Assumption 1 hold, and let the system data satisfy appropriate technical conditions (smoothness, etc.). Then there exists a $\varepsilon^* > 0$ such that for all $\varepsilon \in (0,\varepsilon^*)$ we have

$$x(t) - \bar{x}(t) = O(\varepsilon) \qquad \forall\, t \in [0,T]$$

and

$$z(t) - h(t, \bar{x}(t)) = O(\varepsilon) \qquad \forall\, t \in [t^*, T]$$

This is known as Tikhonov's theorem. See [Khalil, Theorem 11.1] for a precise statement (and a proof).

Geometrically: in the (x, z)-space, consider the surface $z = h(x)$ (assume the system is autonomous). On this surface, the dynamics are defined by $\dot{x} = f(x, h(x))$, $z = h(x)$. Trajectories of the full (x, z)-system approach this surface fast and then stay in an "ε-tube" around it.

The above estimates apply only on a finite interval $[0, T]$. The $O(\varepsilon)$ terms are valid on a compact subset of the state space that contains the state up to time T, and are not in general valid as $t \to \infty$.

We can actually get an approximation on the infinite time interval, if we assume that the reduced system is stable.

Assumption 2: $x = 0$ is an exponentially stable equilibrium of the reduced system (110).

With this additional assumption, the previous result is valid for $T = \infty$. See [Khalil, Theorem 11.2].

The above result does not guarantee exponential stability of the origin. In fact, the origin is in general not an equilibrium point of the full system, for $\varepsilon > 0$; we only know that $(x, \bar{z}) = (0, 0)$ is an equilibrium for $\varepsilon = 0$. Note that $z = h(t, x)$ is not an equilibrium of the z-subsystem unless $\varepsilon = 0$. If we make further assumptions to ensure that $(x, z) = (0, 0)$ is an equilibrium of the full system for $\varepsilon > 0$, then its exponential stability does follow from Assumptions 1 and 2 via Lyapunov analysis. See [Khalil, Theorem 11.4].

10.3 Direct MRAC with unmodeled dynamics

Go back to the system (108). Suppose that we want to solve the MRAC problem for it, with the reference model

$$\dot{y}_m = -a_m y_m + b_m r$$

where $a_m > 0$ and b_m are known and we take the reference signal to be

$$r = \sin \omega t$$

Let us first *ignore the z-dynamics* and follow the direct control approach of Section 6.5.1 where we designed the adaptive controller for the scalar plant (109).

The control law we derived for direct MRAC is

$$u = -\hat{k} y + \hat{l} r \tag{113}$$

and the tuning laws are

$$\dot{\hat{k}} := -\gamma e y, \qquad \dot{\hat{l}} := \gamma e r$$

where $\gamma > 0$.

Now, let's see what happens if we apply the above adaptive controller to the singularly perturbed plant (108). We have to repeat some of the derivations we did earlier for direct MRAC, with minor differences.

First, we need the direct control reparameterization of the y-dynamics in terms of the "nominal" controller gains

$$k = \frac{a + a_m}{b}, \qquad l = \frac{b_m}{b}$$

Compared to what we had in direct MRAC before, there is one additional term which gives the difference between z and u:

$$\begin{aligned}\dot{y} &= -a_m y + b_m r + b(z + ky - lr)\\ &= -a_m y + b_m r + b(u + ky - lr) + b(z - u)\end{aligned}$$

Plugging in the expression (113) for u into this, we get

$$\dot{y} = -a_m y + b_m r + b(-\tilde{k}y + \tilde{l}r) + b(z - u)$$

From this, we see that the output tracking error

$$e := y_m - y$$

satisfies

$$\dot{e} = -a_m e + b\tilde{k}y - b\tilde{l}r - b(z - u)$$

Next, we rewrite the full closed-loop system using the $(e, \tilde{k}, \tilde{l})$ coordinates[15]:

$$\begin{aligned}\begin{pmatrix}\dot{e}\\ \dot{\tilde{k}}\\ \dot{\tilde{l}}\end{pmatrix} &= \underbrace{\begin{pmatrix}-a_m & bA\sin(\omega t + \alpha) & -b\sin\omega t\\ -\gamma A\sin(\omega t + \alpha) & 0 & 0\\ \gamma\sin\omega t & 0 & 0\end{pmatrix}}_{A(t)}\begin{pmatrix}e\\ \tilde{k}\\ \tilde{l}\end{pmatrix} - \begin{pmatrix}b\\ 0\\ 0\end{pmatrix}(z - u)\\ \varepsilon\dot{z} &= -z + u\end{aligned}$$

Define

$$x := \left(e, \tilde{k}, \tilde{l}\right)^T$$

We can write the control law (113) as

$$u = -(\tilde{k} + k)(y_m(t) - e) + (\tilde{l} + l)r(t)$$

[15]Up to terms that can eventually be shown to decay to 0; see footnote 8 in Section 6.5.1.

which takes the form

$$u = h(t, x)$$

where t on the right-hand side takes care of the dependence on $r(t)$ and $y_m(t)$ (which are known functions of time). We obtain the x-dynamics

$$\dot{x} = A(t)x - \begin{pmatrix} b \\ 0 \\ 0 \end{pmatrix} (z - h(t, x))$$

The reduced model is thus

$$\dot{x} = A(t)x$$

and we showed in Section 6.5.1 that it is exponentially stable.

We also need to consider

$$\varepsilon \dot{z} = -z + h(t, x)$$

The system (112) in this case is

$$\frac{d\bar{z}}{d\tau} = -\bar{z} - \cancel{h(t,x)} + \cancel{h(t,x)} = -\bar{z}$$

and this is clearly exponentially stable uniformly over t and x.

Assumptions 1 and 2 of the previous subsection are satisfied, and we conclude that there exists a $\varepsilon^* > 0$ such that for all $\varepsilon \in (0, \varepsilon^*)$ we have

$$x(t) - \bar{x}(t) = O(\varepsilon) \qquad \forall\, t \geq 0$$

In particular, output tracking is achieved up to order ε. This characterizes robustness of our direct MRAC design to fast unmodeled dynamics. (The general result also provides a similar estimate for z, but it's not important for us here.)

Can we try to show that the closed-loop system is exponentially stable for ε small enough?

No, because the origin is not even an equilibrium of the closed-loop system. This is because

$$h(t, 0) = -ky_m(t) + lr(t) \neq 0$$

(See also the last paragraph of the previous subsection.) The effect of unmodeled dynamics does not necessarily diminish with time, the displacement is small but persistent.

Optional exercise: simulate the system and investigate its behavior.

11 Conclusion

So, can we now define what "adaptive control" means?

Last three lectures: final project presentations.