Chapter 1

Introduction

1.1 Historical Overview

Research in adaptive control was motivated by the design of autopilots for highly agile aircraft that need to operate at a wide range of speeds and altitudes, experiencing large parametric variations. In the early 1950s adaptive control was conceived and proposed as a technology for automatically adjusting the controller parameters in the face of changing aircraft dynamics [61, 126]. In [14], that period is called the *brave era* because "there was a very short path from idea to flight test with very little analysis in between." The tragic flight test of the X-15 was the first trial of an adaptive flight control system [164]. It clearly indicated a lack of depth in understanding the robustness properties of adaptive feedback loops.

The initial results in adaptive control were inspired by system identification [115], which led to an architecture consisting of an online parameter estimator combined with automatic control design [16, 81]. Two architectures of adaptive control emerged: the direct method, where only controller parameters were estimated, and the indirect method, where process parameters were estimated and the controller parameters were obtained using some design procedure. To achieve identifiability, it was necessary to introduce a condition of persistency of excitation [15] in order to guarantee that the parameter estimates converge. The relationships between the architectures were clarified in [48].

The progress in systems theory led to fundamental theory for development of stable adaptive control architectures (see [18, 20, 49, 57, 102, 103, 127, 132, 150, 151, 159] and references therein). This was accompanied by several examples, including Rohrs' example, challenging the robustness of adaptive controllers in the presence of unmodeled dynamics, [147]. Although [147] included a rigorous proof of the existence of two infinite-gain operators in the closed-loop adaptive system, the explanation given for the phenomena observed in the simulations, which was based on qualitative considerations, was not complete. A thorough explanation was provided in later papers by Åström [12] and Anderson [5]. Nevertheless, with his example, Rohrs brought up an important point: the available adaptive control algorithms to that date were unable to adjust the bandwidth of the closed-loop system and guarantee its robustness. The results and conclusions of this paper led to an ideological controversy, and other authors started to investigate the robustness and convergence of adaptive controllers.

The works of Ioannou and Kokotović [72–74], Peterson and Narendra [142], Kresselmeier and Narendra [96], and Narendra and Annaswamy [130] deserve special mention. In these papers, the authors analyzed the causes of instability and proposed damping-type modifications of adaptive laws to prevent them. The basic idea of all the modifications was to limit the gain of the adaptation loop and to eliminate its integral action. Examples of these modifications are the σ-modification [74] and the e-modification [130]. All these modifications attempted to provide a solution to the problem of *parameter drift*; however, they did not directly address the *architectural problem* identified by Rohrs. We notice that lack of robustness of adaptive controllers has been analyzed in robust control literature [55]. An incomplete overview of robustness and stability issues of adaptive controllers can be found in [5].

On the other hand, an example presented in [185] demonstrated that the system output can have overly poor transient tracking behavior before ideal asymptotic convergence takes place. In [182], the author proved that it may not be possible to optimize $\mathcal{L}_2$ and $\mathcal{L}_\infty$ performance simultaneously by using a constant adaptation rate. Following these results, modifications of adaptive controllers were proposed in [43, 163] that render the tracking error arbitrarily small in terms of both mean-square and $\mathcal{L}_\infty$-bounds. Further, it was shown in [42] that the modifications proposed in [43, 163] could be derived as a linear feedback of the tracking error, and the improved performance was obtained only due to a nonadaptive high-gain feedback. In [159], a composite adaptive controller was proposed, which suggests a new adaptation law using both tracking error and prediction error and leads to less oscillatory behavior in the presence of high adaptation gains as compared to model reference adaptive control (MRAC). In [125], a high-gain switching MRAC technique was introduced to achieve arbitrary good transient tracking performance under a relaxed set of assumptions as compared to MRAC, and the results were shown to be of existence type only. In [131], a multiple model switching scheme was proposed to improve the transient performance of adaptive controllers. In [10], it was shown that an arbitrarily close transient bound can be achieved by enforcing a parameter-dependent persistent excitation condition. In [101], computable $\mathcal{L}_2$- and $\mathcal{L}_\infty$-bounds for the output tracking error signals were obtained for a special class of adaptive controllers using backstepping. The underlying linear nonadaptive controller possesses a parametric robustness property. However, for a large parametric uncertainty it requires high-gain feedback. In [136], dynamic certainty equivalent controllers with unnormalized estimators were used for adaptation, which permit derivation of a uniform upper bound for the $\mathcal{L}_2$-norm of the tracking error in terms of the initial parameter estimation error. In the presence of sufficiently small initial conditions, the author proved that the $\mathcal{L}_\infty$-norm of the tracking error is upper bounded by the $\mathcal{L}_\infty$-norm of the reference input. In [9, 50, 137], a differential game theoretic type $\mathcal{H}_\infty$ approach was investigated for achieving arbitrarily close disturbance attenuation for tracking performance, albeit at the price of increased control effort. In [187], a new certainty-equivalence-based adaptive controller was presented using a backstepping-type controller with a normalized adaptive law to achieve asymptotic stability and guarantee performance bounds comparable with the tuning functions scheme, without the use of higher-order nonlinearities. References [128, 129] developed the supervisory control approach that defines a fast switching scheme between candidate controllers leading to guaranteed performance bounds. However, robustness of these schemes to unmodeled dynamics appears to be limited by the frequency of switching [6, 66].

As compared to the linear systems theory, several important aspects of the transient performance analysis seem to be missing in these efforts. First, the bounds are computed

for tracking errors only, not for control signals. Although the latter can be deduced from the former, it is straightforward to verify that the ability to adjust the former may not extend to the latter in case of nonlinear control laws. Second, since the purpose of adaptive control is to ensure stable performance in the presence of modeling uncertainties, one needs to ensure that (admissible) changes in reference commands and system dynamics due to possible faults or unexpected uncertainties do not lead to unacceptable transient deviations or oscillatory control signals, implying that a *retuning of adaptation parameters* is required. Finally, one needs to ensure that the modifications or solutions, suggested for performance improvement of adaptive controllers, are not achieved via high-gain feedback.

In brief summary, the development of the theory of adaptive control over the years has taken rather the trend of defining a *larger and larger class* of systems, for which a Lyapunov proof can be done for asymptotic stability. *At which location should the uncertainty appear, what should be the degree of mismatch, how should the adaptive law be modified, etc., to get a negative definite (semidefinite) derivative of the associated candidate Lyapunov function for* a new class *of systems?* These questions or one of them is present in almost every paper addressing the next stage of development in the theory of adaptive control. Significant efforts have been reported on relaxation of the matching conditions by extending the backstepping-design approach to a broader class of systems, including strict-parametric feedback and feedforward systems [9, 45, 97, 100, 137, 138], analysis of robustness of these schemes to unmodeled dynamics [4,8,70,71,77,134], extensions to output feedback with an objective to achieve global or semiglobal output feedback stabilization [76,84,98,99,119,120], extension to systems with time-varying parameters [68, 122, 135, 186], relaxation of the relative degree [strictly positive real (SPR)] requirement via input-filtered transformations [118, 121], extension to nonminimum phase systems [69, 75], etc.

These fundamental results provide sufficient conditions on the bounds of uncertainties and initial conditions, which would guarantee that with the given adaptive feedback architecture, the signals in the feedback loop remain bounded. Though very important, when dealing with practical applications, boundedness, ultimate boundedness, or even asymptotic convergence are weak properties for nonlinear (adaptive) feedback systems. On one hand, unmodeled dynamics, latencies, and noise require precise quantification of the robustness and the stability margins of the underlying feedback loop. On the other hand, performance requirements in real applications necessitate a *predictable* response for the closed-loop system, dependent upon the changes in system dynamics. In adaptive control, the nature of the adaptation process plays a central role in both robustness and performance. Ideally, one would like *adaptation* to *correctly* respond to all the changes in initial conditions, reference inputs, and uncertainties by *quickly identifying* a set of control parameters that would provide a satisfactory system response. This, of course, demands *fast estimation schemes with high adaptation rates* and, as a consequence, leads to the fundamental question of determining the upper bound on the adaptation rate that would *not* result in poor robustness characteristics. We notice that the results available in the literature consistently limited the rate of variation of uncertainties, by providing examples of destabilization due to fast adaptation [75, p. 549], while the transient performance analysis was continually reduced to persistency of excitation-type assumptions, which, besides being a highly undesirable phenomenon, cannot be verified a priori. The lack of analytical quantification of the relationship between the rate of adaptation, the transient response, and the robustness margins led to *gain-scheduled designs of adaptive controllers*, examples of which are the successful flight tests of the late 1990s by the Air Force and Boeing [175, 176]. The flight tests relied

on intensive Monte Carlo analysis for determination of the best rate of adaptation for various flight conditions. It was apparent that fast adaptation was leading to high frequencies in control signals and increased sensitivity to time delays. The fundamental question was thus reduced to determining an *architecture*, which would allow for *fast adaptation without losing robustness*. It was clearly understood that such an architecture can reduce the amount of gain scheduling, and possibly eliminate it, as *fast adaptation—in the presence of guaranteed robustness*—should be able to compensate for the negative effects of rapid variation of uncertainties on the system response.

The $\mathcal{L}_1$ adaptive control theory addressed precisely this question by setting an architecture in place for which adaptation is *decoupled* from robustness. The speed of adaptation in these architectures is limited only by the available hardware, while robustness is resolved via conventional methods from classical and robust control. The architectures of $\mathcal{L}_1$ adaptive control theory have *guaranteed transient performance* and *guaranteed robustness* in the presence of *fast adaptation*, without introducing or enforcing persistence of excitation, without any gain scheduling in the controller parameters, and without resorting to high-gain feedback. With $\mathcal{L}_1$ adaptive controller in the feedback loop, the response of the closed-loop system can be predicted a priori, thus significantly reducing the amount of Monte Carlo analysis required for verification and validation of such systems. These features of $\mathcal{L}_1$ adaptive control theory were verified—*consistently with the theory*—in a large number of flight tests and in mid- to high-fidelity simulation environments [19, 35, 36, 46, 51, 58, 59, 67, 82, 83, 88, 94, 104–106, 110, 124, 140, 141, 170, 172].

To facilitate the development of $\mathcal{L}_1$ adaptive control theory, in the next section we introduce two equivalent architectures of MRAC, which lead to the same error dynamics from the same initial conditions. We later use one of these structures as a basis for development of the main results in this book.

1.2 Two Different Architectures of Adaptive Control

In this section we present two different, but equivalent, architectures of adaptive control. Although their implementation is different, they both lead to the same error dynamics from the same initial conditions. The difference in their implementation principle is the key to the development of $\mathcal{L}_1$ adaptive control architectures in this book.

1.2.1 Direct MRAC

Let the system dynamics propagate according to the following differential equation:

$$\begin{aligned} \dot{x}(t) &= A_m x(t) + b\left(u(t) + k_x^\top x(t)\right), \quad x(0) = x_0, \\ y(t) &= c^\top x(t), \end{aligned} \tag{1.1}$$

where $x(t) \in \mathbb{R}^n$ is the state of the system (measured), $A_m \in \mathbb{R}^{n\times n}$ is a known Hurwitz matrix that defines the desired dynamics for the closed-loop system, $b, c \in \mathbb{R}^n$ are known constant vectors, $k_x \in \mathbb{R}^n$ is a vector of unknown constant parameters, $u(t) \in \mathbb{R}$ is the control input, and $y(t) \in \mathbb{R}$ is the regulated output. Given a uniformly bounded piecewise-continuous

reference input $r(t) \in \mathbb{R}$, the objective is to define an adaptive feedback signal $u(t)$ such that $y(t)$ tracks $r(t)$ with desired specifications, while all the signals remain bounded.

The MRAC architecture proceeds by considering the nominal controller

$$u_{\text{nom}}(t) = -k_x^\top x(t) + k_g r(t), \tag{1.2}$$

where

$$k_g \triangleq \frac{1}{c^\top A_m^{-1} b}. \tag{1.3}$$

This nominal controller assumes perfect cancelation of the uncertainties in (1.1) and leads to the desired (ideal) reference system

$$\begin{aligned} \dot{x}_m(t) &= A_m x_m(t) + b k_g r(t), \quad x_m(0) = x_0, \\ y_m(t) &= c^\top x_m(t), \end{aligned} \tag{1.4}$$

where $x_m(t) \in \mathbb{R}^n$ is the state of the reference model. The choice of k_g according to (1.3) ensures that $y_m(t)$ tracks step reference inputs with zero steady-state error.

The direct model reference adaptive controller is given by

$$u(t) = -\hat{k}_x^\top(t) x(t) + k_g r(t), \tag{1.5}$$

where $\hat{k}_x(t) \in \mathbb{R}^n$ is the estimate of k_x. Substituting (1.5) into (1.1) yields the closed-loop system dynamics

$$\begin{aligned} \dot{x}(t) &= (A_m - b\tilde{k}_x^\top(t)) x(t) + b k_g r(t), \quad x(0) = x_0, \\ y(t) &= c^\top x(t), \end{aligned}$$

where $\tilde{k}_x(t) \triangleq \hat{k}_x(t) - k_x$ denotes the parametric estimation error.

Letting $e(t) \triangleq x_m(t) - x(t)$ be the tracking error signal, the tracking error dynamics can be written as

$$\dot{e}(t) = A_m e(t) + b\tilde{k}_x^\top(t) x(t), \quad e(0) = 0. \tag{1.6}$$

The update law for the parametric estimate is given by

$$\dot{\hat{k}}_x(t) \quad = \quad -\Gamma x(t) e^\top(t) P b, \quad \hat{k}_x(0) = k_{x0}, \tag{1.7}$$

where $\Gamma \in \mathbb{R}^+$ is the adaptation gain and $P = P^\top > 0$ solves the algebraic Lyapunov equation

$$A_m^\top P + P A_m = -Q$$

for arbitrary $Q = Q^\top > 0$. The block diagram of the closed-loop system is given in Figure 1.1.

Consider the following Lyapunov function candidate:

$$V(e(t), \tilde{k}_x(t)) = e^\top(t) P e(t) + \frac{1}{\Gamma} \tilde{k}_x^\top(t) \tilde{k}_x(t). \tag{1.8}$$

Its time derivative along the system trajectories (1.6)–(1.7) is given by

$$\begin{aligned}\dot{V}(t) &= -e^{\top}(t)Qe(t) + 2e^{\top}(t)Pb\tilde{k}_x^{\top}(t)x(t) + \frac{2}{\Gamma}\tilde{k}_x^{\top}(t)\dot{\tilde{k}}_x(t)\\ &= -e^{\top}(t)Qe(t) + 2\tilde{k}_x^{\top}(t)\left(\frac{1}{\Gamma}\dot{\tilde{k}}_x(t) + x(t)e^{\top}(t)Pb\right)\\ &= -e^{\top}(t)Qe(t) \leq 0\,.\end{aligned}$$

Hence, the equilibrium of (1.6)–(1.7) is Lyapunov stable, i.e., the signals $e(t)$, $\tilde{k}_x(t)$ are bounded. Since $x(t) = x_m(t) - e(t)$, and $x_m(t)$ is the state of a stable reference model, then $x(t)$ is bounded. To show that the tracking error converges asymptotically to zero, we compute the second derivative of $V(e(t), \tilde{k}_x(t))$ as

$$\ddot{V}(t) = -2e^{\top}(t)Q\dot{e}(t)\,.$$

It follows from (1.6) that $\dot{e}(t)$ is uniformly bounded, and hence $\ddot{V}(t)$ is bounded, implying that $\dot{V}(t)$ is uniformly continuous. Application of Barbalat's lemma (see A.6.1) yields

$$\lim_{t\to\infty} \dot{V}(t) = 0\,,$$

which consequently proves that $e(t) \to 0$ as $t \to \infty$. Thus, $x(t)$ asymptotically converges to $x_m(t)$. This in turn implies that $y(t) = c^{\top}x(t)$ asymptotically converges to $y_m(t) = c^{\top}x_m(t)$, which follows $r(t)$ with desired specifications.

Notice that asymptotic convergence of parametric estimation errors to zero is not guaranteed. The parametric estimation errors are guaranteed only to stay bounded.

1.2.2 Direct MRAC with State Predictor

Next, we consider a reparameterization of the above architecture using a state predictor (or identifier), given by

$$\begin{aligned}&\dot{\hat{x}}(t) = A_m\hat{x}(t) + b(u(t) + \hat{k}_x^{\top}(t)x(t)), \quad \hat{x}(0) = x_0\,,\\ &\hat{y}(t) = c^{\top}\hat{x}(t)\,,\end{aligned} \tag{1.9}$$

where $\hat{x}(t) \in \mathbb{R}^n$ is the state of the predictor. The system in (1.9) replicates the system structure from (1.1) with the unknown parameter k_x replaced by its estimate $\hat{k}_x(t)$. By subtracting (1.1) from (1.9), we obtain the *prediction error dynamics* (or identification error dynamics), independent of the control choice,

$$\dot{\tilde{x}}(t) = A_m\tilde{x}(t) + b\tilde{k}_x^{\top}(t)x(t), \quad \tilde{x}(0) = 0\,,$$

where $\tilde{x}(t) \triangleq \hat{x}(t) - x(t)$ and $\tilde{k}_x(t) \triangleq \hat{k}_x(t) - k_x$. Notice that these error dynamics are identical to the error dynamics in (1.6).

Next, let the adaptive law for $\hat{k}_x(t)$ be given as

$$\dot{\hat{k}}_x(t) = -\Gamma x(t)\tilde{x}^{\top}(t)Pb, \quad \hat{k}_x(0) = k_{x0}\,, \tag{1.10}$$

where $\Gamma \in \mathbb{R}^+$ is the adaptation rate and $A_m^\top P + P A_m = -Q$, $Q = Q^\top > 0$. This adaptive law is similar to (1.7) in its structure, except that the tracking error $e(t)$ is replaced by the prediction error $\tilde{x}(t)$. The choice of the Lyapunov function candidate

$$V(\tilde{x}(t), \tilde{k}_x(t)) = \tilde{x}^\top(t) P \tilde{x}(t) + \frac{1}{\Gamma} \tilde{k}_x^\top(t) \tilde{k}_x(t)$$

leads to

$$\dot{V}(t) = -\tilde{x}^\top(t) Q \tilde{x} \leq 0,$$

implying that the errors $\tilde{x}(t)$ and $\tilde{k}_x(t)$ are uniformly bounded. Notice, however, that without introducing the feedback signal $u(t)$ one cannot apply Barbalat's lemma to conclude asymptotic convergence of $\tilde{x}(t)$ to zero. Both $x(t)$ and $\hat{x}(t)$ can diverge at the same rate, keeping $\tilde{x}(t)$ uniformly bounded.

If we use (1.5) in (1.9), with account of (1.10), the closed-loop state predictor replicates the bounded reference system of (1.4):

$$\begin{aligned} \dot{\hat{x}}(t) &= A_m \hat{x}(t) + b k_g r(t), \quad \hat{x}(0) = x_0, \\ \hat{y}(t) &= c^\top \hat{x}(t). \end{aligned}$$

Hence, Barbalat's lemma can be invoked to conclude that $\tilde{x}(t) \to 0$ as $t \to \infty$. The block diagram of the closed-loop system with the predictor is given in Figure 1.2.

Figures 1.1 and 1.2 illustrate the fundamental difference between the direct MRAC and the predictor-based adaptation. In Figure 1.2, the control signal is provided as input to both systems, the system and the predictor, while in Figure 1.1 the control signal serves only as input to the system. This feature is the key to the development of $\mathcal{L}_1$ adaptive control architectures with quantifiable performance bounds.

1.2.3 Tuning Challenges

From the above Lyapunov analysis, we notice that the tracking error can be upper bounded in the following way:

$$\|e(t)\| (= \|\tilde{x}(t)\|) \leq \sqrt{\frac{V(t)}{\lambda_{\min}(P)}} \leq \sqrt{\frac{V(0)}{\lambda_{\min}(P)}} = \frac{\|\tilde{k}_x(0)\|}{\sqrt{\lambda_{\min}(P)\Gamma}}, \quad \forall\, t \geq 0.$$

This bound shows that the tracking error can be arbitrarily reduced for all $t \geq 0$ (including the transient phase) by increasing the adaptation gain Γ [100]. However, from the control law in (1.5) and the adaptive laws in (1.7) and (1.10), it follows that large adaptive gains result in high-gain feedback control, which manifests itself in high-frequency oscillations in the control signal and reduced tolerance to time delays. Moreover, applications requiring identification schemes with time scales comparable with those of the closed-loop dynamics appear to be extremely challenging due to undesirable interactions of the two processes [5]. Due to the lack of systematic design guidelines to select an adequate adaptation gain, tuning of such applications is being commonly resolved by either computationally expensive Monte Carlo simulations or trial-and-error methods following some empirical guidelines or engineering intuition. As a consequence, proper tuning of MRAC architectures (or their equivalent state-predictor-based reparameterizations) represents a major challenge and has largely remained an open question in the literature.

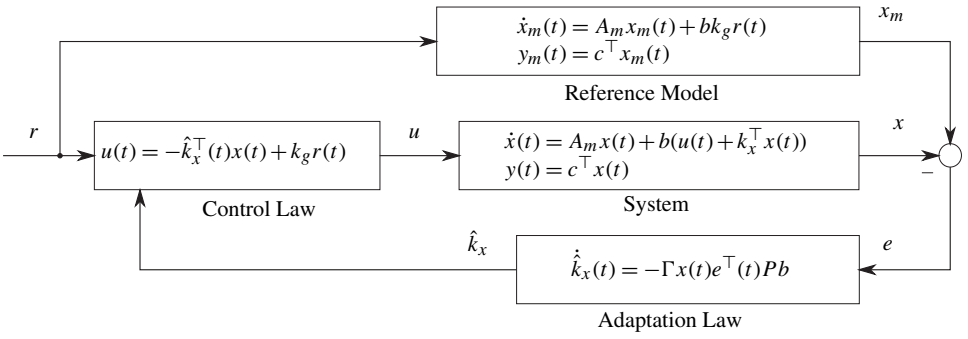

Figure 1.1: Closed-loop direct MRAC architecture.

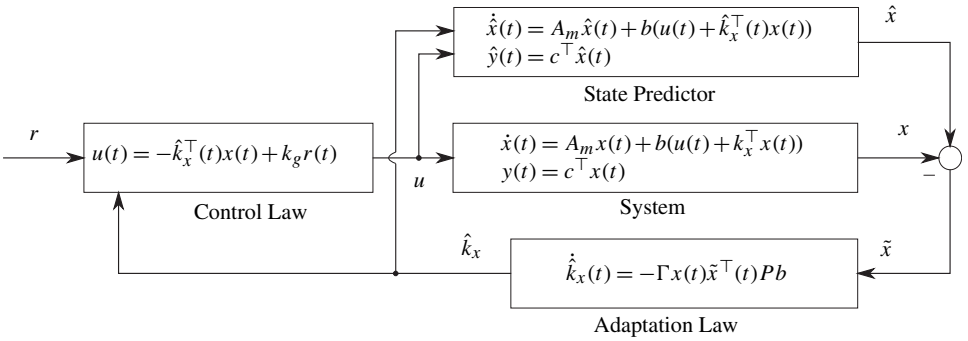

Figure 1.2: Closed-loop MRAC architecture with state predictor.

1.3 Saving the Time-Delay Margin

Next we will introduce the key ideas of the $\mathcal{L}_1$ adaptive controller, which enables *fast adaptation* with *guaranteed robustness*. We will start with a simple stable scalar system with constant disturbance, which can be analyzed by resorting to tools from classical control. We notice that in this case, MRAC reduces to a linear (model-following) integral controller. Since the closed-loop system remains linear, we use the Nyquist criterion to analyze stability and robustness of this system. Taking advantage of its linear structure, we present (i) some of the benefits of $\mathcal{L}_1$ adaptive control architectures and (ii) different concepts and tools that will be used throughout the book. In particular, we will show that *fast adaptation of $\mathcal{L}_1$ adaptive control architectures is beneficial for robustness*. We will also derive the *uniform* performance bounds of the $\mathcal{L}_1$ adaptive controller, for both the state and the control signal, and show the role of the bandwidth-limited filter of the $\mathcal{L}_1$ architecture in obtaining these uniform bounds.

Toward that end, consider the scalar system

$$\dot{x}(t) = -x(t) + \theta + u(t), \quad x(0) = x_0, \tag{1.11}$$

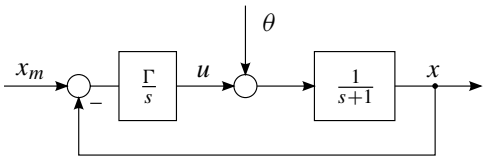

Figure 1.3: Closed-loop system with MRAC-type integral controller.

where θ is the unknown constant to be rejected by the control input $u(t)$. Let the objective be stabilization of the origin. For this system, the MRAC architecture described in (1.4) and (1.5) reduces to an integral controller of the structure

$$u(t) = -\hat{\theta}(t), \tag{1.12}$$

where $\hat{\theta}(t)$ is the estimate of θ, given by

$$\dot{\hat{\theta}}(t) = -\Gamma(x_m(t) - x(t)), \quad \theta(0) = \theta_0, \quad \Gamma > 0, \tag{1.13}$$

and $x_m(t)$ is the reference signal, generated by the system

$$\dot{x}_m(t) = -x_m(t), \quad x_m(0) = x_0.$$

We notice that this reference system is obtained from the original system (1.11) by substitution of the ideal nominal controller $u_{\text{nom}}(t) = -\theta$ into it, thus assuming perfect cancelation of the uncertain parameter θ in the system (1.11). The block diagram of the closed-loop system is shown in Figure 1.3.

The loop transfer function of this system (with negative feedback) is

$$L_1(s) = \frac{\Gamma}{s(s+1)}. \tag{1.14}$$

Because the closed-loop system remains linear time-invariant (LTI), one can use standard tools from classical control to analyze the stability margins of this system. The two most commonly used stability margins are the gain and the phase margin. From Figure 1.4(a), it is obvious that the Nyquist plot of $L_1(s)$ never crosses the negative part of the real line; therefore, the closed-loop system has infinite gain margin ($g_m = \infty$). The gain crossover frequency ω_{gc} can be computed from

$$|L_1(j\omega_{gc})| = \frac{\Gamma}{\omega_{gc}\sqrt{\omega_{gc}^2 + 1}} = 1,$$

which leads to the phase margin

$$\phi_m = \pi + \angle L_1(j\omega_{gc}) = \arctan\left(\frac{1}{\omega_{gc}}\right).$$

Careful analysis indicates that increasing Γ leads to higher gain crossover frequency and consequently reduces the phase margin. The reduction of phase margin with large Γ can

also be observed in Figure 1.4(a). So, if increasing Γ improves the tracking performance for all $t \geq 0$, including the *transient phase*, then it obviously hurts the robustness (or relative stability) of the closed-loop system. Thus, the adaptation rate Γ is the key to the trade-off between performance and robustness. Since tracking and robustness cannot be achieved simultaneously, there is nothing surprising about this, but we would like to explore if the architecture can be modified so that the trade-off between tracking and robustness is resolved differently and the adaptation gain Γ can be safely increased for transient performance improvement without hurting the robustness of the closed-loop system (see Section 1.2.3).

To obtain the $\mathcal{L}_1$ adaptive controller for this system, the controller in (1.12)–(1.13) will be modified in two ways. First, we introduce the state predictor,

$$\dot{\hat{x}}(t) = -\hat{x}(t) + \hat{\theta}(t) + u(t), \quad \hat{x}(0) = x_0,$$

which leads to the following prediction error dynamics, independent of the control choice:

$$\dot{\tilde{x}}(t) = -\tilde{x}(t) + \tilde{\theta}(t), \quad \tilde{x}(0) = 0, \tag{1.15}$$

where $\tilde{x}(t) \triangleq \hat{x}(t) - x(t)$ and $\tilde{\theta}(t) \triangleq \hat{\theta}(t) - \theta$. The parametric estimate, given by (1.13), is thus replaced by

$$\dot{\hat{\theta}}(t) = -\Gamma \tilde{x}(t), \quad \theta(0) = \theta_0, \quad \Gamma > 0.$$

Next, instead of choosing the adaptive controller as $u(t) = -\hat{\theta}(t)$, we use a low-pass filtered version of $\hat{\theta}(t)$,

$$u(s) = -C(s)\hat{\theta}(s), \tag{1.16}$$

where $u(s)$ and $\hat{\theta}(s)$ are the Laplace transforms of $u(t)$ and $\hat{\theta}(t)$, respectively, and $C(s)$ is a bounded-input bounded-output (BIBO) stable strictly proper transfer function subject to $C(0) = 1$ with zero initialization for its state-space realization. The block diagram of this system is given in Figure 1.5. In the foregoing analysis, we further consider a first-order low-pass filter

$$C(s) = \frac{\omega_c}{s + \omega_c}; \tag{1.17}$$

however, similar results can be obtained using more complex filters. The loop transfer function of this system (with negative feedback) is

$$L_2(s) = \frac{\Gamma C(s)}{s(s+1) + \Gamma(1 - C(s))}. \tag{1.18}$$

Notice that in the absence of the filter, i.e., with $C(s) = 1$, the controller in (1.16) reduces to the MRAC-type integral controller introduced earlier, and (1.18) reduces to (1.14), that is, $L_2(s) = L_1(s)$.

Although (1.18) has a more complex structure than (1.14), the Nyquist plot in Figure 1.4(b) shows that the phase and the gain margins of the $\mathcal{L}_1$ controller are not significantly affected by large values of Γ. The effect of the adaptive gain on the robustness margins of the two closed-loop systems is clearly presented in Figure 1.6. The figure shows that, while the phase margin of the MRAC-type integral controller vanishes as one increases the adaptation gain Γ, the $\mathcal{L}_1$ adaptive controller has a guaranteed bounded-away-from-zero phase and gain margins in the presence of *fast adaptation*.

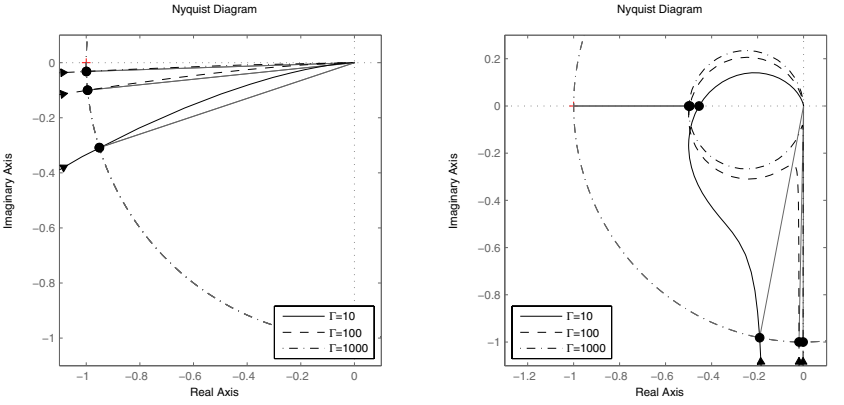

(a) Integral controller (b) $\mathcal{L}_1$ controller

Figure 1.4: Nyquist plots for the loop transfer functions.

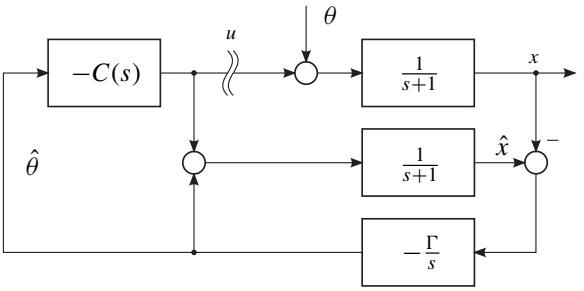

Figure 1.5: Closed-loop system with $\mathcal{L}_1$ adaptive controller.

Further, notice that as $\Gamma \to \infty$, the expression in (1.18) leads to the following limiting loop transfer function:

$$L_{2l}(s) = \frac{C(s)}{1 - C(s)} = \frac{\omega_c}{s}. \tag{1.19}$$

This loop transfer function has an infinite gain margin ($g_m = \infty$) and a phase margin of $\phi_m = \pi/2$. However, from Figure 1.6(a), we notice that the gain margin is always finite and actually converges to $g_m = 6.02$ dB with the increase of Γ. We note that the (high-frequency) dynamics of the adaptation loop do not appear in the limiting loop transfer function in (1.19). Then, since the phase crossover frequency tends to infinity as the adaptation gain Γ increases, this limiting loop transfer function cannot be used to analyze the gain margin of the closed-loop system with the $\mathcal{L}_1$ adaptive controller. However, the gain crossover frequency stays in the low-frequency range, where the limiting loop transfer function in (1.19) is a good approximation of the actual loop transfer function in (1.18). Consequently, the limiting loop transfer function can be used to analyze the phase margin of the closed-loop adaptive system.

One can equivalently measure the robustness of the system by computing its *time-delay margin*, which is defined as the amount of time delay that brings the system to the

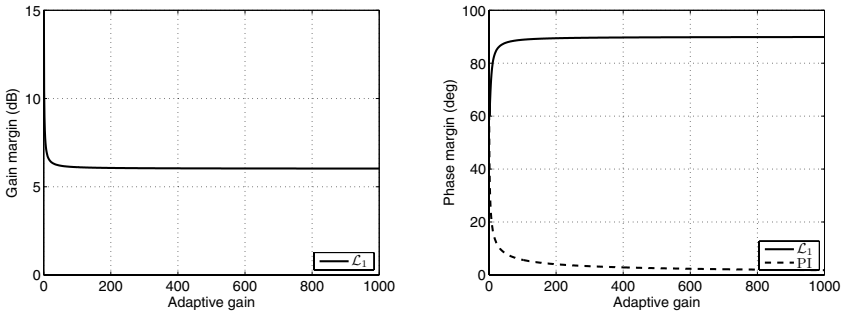

(a) gain margin (b) phase margin

Figure 1.6: Effect of high adaptation gain on the stability margins.

verge of instability. The additional phase lag in the system due to a time delay τ is given by $\phi_\tau(\omega) = \angle\left(e^{-\tau j\omega}\right) = -\tau\omega$. Recalling the definition of the phase margin, one can compute the time-delay margin $\mathcal{T}$ as the amount of delay introduced in the system that reduces the phase margin to zero:

$$\phi_m = \mathcal{T}\omega_{gc} \quad \Rightarrow \quad \mathcal{T} = \frac{\phi_m}{\omega_{gc}}.$$

From (1.19), it follows that $\omega_{gc} = \omega_c$, which further implies that the $\mathcal{L}_1$ adaptive controller has the following time-delay margin as $\Gamma \to \infty$:

$$\mathcal{T} = \frac{\phi_m}{\omega_{gc}} = \frac{\pi}{2\omega_c}.$$

Hence, we observe that the $\mathcal{L}_1$ adaptive controller, defined by (1.16), retains guaranteed robustness in the presence of large values of Γ, while the MRAC-type integral controller obviously loses its phase margin in the presence of fast adaptation.

1.4 Uniformly Bounded Control Signal

Next, we analyze a key property of the $\mathcal{L}_1$ adaptive controller, which is inherently related to the robustness features discussed above. We start by considering the following closed-loop structure:

$$\begin{aligned} x_{\text{ref}}(s) &= \frac{1}{s+1}\left(\frac{\theta}{s} + u_{\text{ref}}(s)\right) + \frac{x_0}{s+1}, \\ u_{\text{ref}}(s) &= -C(s)\frac{\theta}{s}. \end{aligned} \tag{1.20}$$

This system is constructed from (1.11) and (1.16) by using $\theta/s = \mathfrak{L}(\theta)$ instead of $\hat{\theta}(s)$ in (1.16) and, hence, represents a closed-loop architecture using the *ideal* nonadaptive version of the $\mathcal{L}_1$ controller. We will refer to this system as a *(closed-loop) reference system*, as it is with respect to this system that we are able to compute uniform performance bounds.

Notice that the reference controller $u_{\text{ref}}(s) = -C(s)\frac{\theta}{s}$, as compared to the nominal controller $u_{\text{nom}}(s) = -\frac{\theta}{s}$ of MRAC, assumes only partial cancelation of uncertainties, i.e., it compensates only for the uncertainties within the bandwidth of $C(s)$. This reference system defines the *best achievable performance* with the $\mathcal{L}_1$ adaptive architecture. The response of this closed-loop reference system can be written as

$$x_{\text{ref}}(s) = \frac{1}{s(s+1)}(1-C(s))\theta + \frac{x_0}{s+1}.$$

Similarly, the response of the system in (1.11) with the $\mathcal{L}_1$ controller in (1.16) takes the form (in the frequency domain)

$$x(s) = \frac{1}{s+1}\left(\frac{\theta}{s} - C(s)\hat{\theta}(s)\right) + \frac{x_0}{s+1} = \frac{1}{s+1}\left((1-C(s))\frac{\theta}{s} - C(s)\tilde{\theta}(s)\right) + \frac{x_0}{s+1},$$

where $\hat{\theta}(s)$ and $\tilde{\theta}(s)$ are the Laplace transforms of $\hat{\theta}(t)$ and $\tilde{\theta}(t)$, respectively. Notice that

$$x_{\text{ref}}(s) - x(s) = \frac{1}{s+1}C(s)\tilde{\theta}(s). \tag{1.21}$$

Also, it follows from (1.15) that

$$\tilde{x}(s) = \frac{1}{s+1}\tilde{\theta}(s), \tag{1.22}$$

which allows for rewriting (1.21) as

$$x_{\text{ref}}(s) - x(s) = C(s)\tilde{x}(s). \tag{1.23}$$

Moreover, notice that

$$\tilde{\theta}(s) = \hat{\theta}(s) - \theta/s = -\Gamma\tilde{x}(s)/s + \hat{\theta}_0/s - \theta/s.$$

Substituting the above expression into (1.22) and solving for $\tilde{x}(s)$ leads to

$$\tilde{x}(s) = -\frac{1}{s^2+s+\Gamma}(\theta - \hat{\theta}_0).$$

We can now take the inverse Laplace transform of $\tilde{x}(s)$ for $\Gamma > 1/4$ to obtain

$$\tilde{x}(t) = -\frac{\theta - \hat{\theta}_0}{\sqrt{\Gamma - 1/4}}\mathrm{e}^{-\frac{1}{2}t}\sin(\sqrt{\Gamma - 1/4}\,t). \tag{1.24}$$

This expression yields the following uniform upper bound on the prediction error:

$$|\tilde{x}(t)| \leq \frac{|\theta - \hat{\theta}_0|}{\sqrt{\Gamma - 1/4}}, \quad \forall\, t \geq 0.$$

Letting $\gamma_0 \triangleq |\theta - \hat{\theta}_0|/\sqrt{\Gamma - 1/4}$, we can write

$$|\tilde{x}(t)| \leq \gamma_0, \qquad \forall\, t \geq 0.$$

Notice that $\lim_{\Gamma\to\infty}\gamma_0 = 0$. Also notice from (1.24) that $\lim_{t\to\infty}\tilde{x}(t) = 0$.

From (1.23) we can also derive the following uniform upper bound:

$$|x_{\rm ref}(t)-x(t)| = \left|\int_0^t h_c(\tau)\tilde{x}(t-\tau)d\tau\right| \le \int_0^t |h_c(\tau)\tilde{x}(t-\tau)|d\tau$$
$$\le \gamma_0 \int_0^t |h_c(\tau)|d\tau \le \gamma_0 \int_0^\infty |h_c(\tau)|d\tau\,, \quad \forall\, t \ge 0\,,$$

where $h_c(t)$ is the impulse response of $C(s)$. In particular, for the $C(s)$ in (1.17), the impulse response can be explicitly computed, leading to the following uniform upper bound:

$$|x_{\rm ref}(t)-x(t)| \le \gamma_0\,, \quad \forall\, t \ge 0\,.$$

This implies that the error between the closed-loop system with the $\mathcal{L}_1$ adaptive controller and the closed-loop reference system, which uses the reference controller, can be uniformly bounded by a constant inverse proportional to the square root of the rate of adaptation.

Similarly, using (1.16), (1.20), and (1.22), we can derive

$$u_{\rm ref}(s)-u(s) = C(s)\tilde{\theta}(s) = C(s)(s+1)\tilde{x}(s)\,. \tag{1.25}$$

Denoting $H_u(s) \triangleq C(s)(s+1)$ and letting $h_u(t)$ be the impulse response for $H_u(s)$, we obtain the following upper bound:

$$|u_{\rm ref}(t)-u(t)| \le \gamma_0 \int_0^\infty |h_u(\tau)|d\tau\,, \quad \forall\, t \ge 0\,. \tag{1.26}$$

Because $C(s)$ is strictly proper and BIBO stable, $H_u(s) \triangleq C(s)(s+1)$ is proper and BIBO stable, and hence it has uniformly bounded impulse response, that is $\int_0^\infty |h_u(\tau)|d\tau < \infty$. Further, since $\lim_{\Gamma\to\infty}\gamma_0 = 0$, we can conclude from (1.26) that the time history of the $\mathcal{L}_1$ adaptive controller can be rendered arbitrarily close to the one of the reference controller for all $t \ge 0$ by increasing the rate of adaptation Γ.

Notice that without the low-pass filter, i.e., with $C(s) = 1$, equation (1.25) reduces to

$$u_{\rm ref}(s)-u(s) = (s+1)\tilde{x}(s)\,.$$

From this expression, it is obvious that the transfer function from $\tilde{x}(t)$ to $u_{\rm ref}(t)-u(t)$ is improper, and hence, in the absence of the filter $C(s)$, one cannot *uniformly* upper bound $|u_{\rm ref}(t)-u(t)|$ as we did in (1.26).

This simple analysis illustrates the role of $C(s)$ toward obtaining *a uniform performance bound* for the control signal of the $\mathcal{L}_1$ adaptive control architecture, as compared to its nonadaptive version (which is uniformly bounded by definition). We further notice that this uniform bound is inverse proportional to the square root of the rate of adaptation, similar to the tracking error. Thus, both performance bounds can be systematically improved by increasing the rate of adaptation. The remaining issue is the design of the low-pass linear filter $C(s)$ to ensure that the reference system in (1.20) achieves desired performance specifications in the presence of unknown θ.

In the remainder of this book, we will show that, similar to this simple example and with appropriate extension of the above described concepts to nonlinear closed-loop adaptive systems, the $\mathcal{L}_1$ adaptive control theory shifts the tuning issue from determining the rate of the *nonlinear* gradient minimization scheme to the design of a *linear* strictly

proper and stable filter, implying that the trade-off between performance and robustness of the closed-loop adaptive system can be systematically addressed using well-established tools from classical and robust control.

Finally, we note that the uniform bounds for the system's state and control signals are expressed in terms of the impulse response of proper BIBO-stable transfer functions, which correspond to the $\mathcal{L}_1$-norms of the underlying systems. Consequently, the corresponding control architectures are referred to as $\mathcal{L}_1$ *adaptive controllers*.