Statistical Learning Control of Uncertain Systems: It is better than it seems

V. Koltchinskii, C.T. Abdallah*, M. Ariola, P. Dorato, D. Panchenko

V. Koltchinskii is with the Department of Mathematics and Statistics, University of New Mexico, Albuquerque, NM 87131, USA. E-mail: vlad@math.unm.edu

***Corresponding Author:** C.T. Abdallah is with the Department of EECE, University of New Mexico, Albuquerque, NM 87131, USA. E-mail: chaouki@eece.unm.edu

M. Ariola is with the Dipartimento di Informatica e Sistemistica, Universitá degli Studi di Napoli Federico II, Napoli, Italy. E-mail: ariola@disna.dis.unina.it

P. Dorato is with the Department of EECE, University of New Mexico, Albuquerque, NM 87131, USA. E-mail: peter@eece.unm.edu

D. Panchenko is with the Department of Mathematics and Statistics, University of New Mexico, Albuquerque, NM 87131, USA. E-mail: panchenk@math.unm.edu

The research of V. Koltchinskii is partially supported by NSA Grant MDA904-99-1-0031. The research of C.T. Abdallah is partially supported by Boeing Computer Services Grant 3-48181, and by NSF INT-9818312. The research of M. Ariola is partially supported by the MURST.

Abstract

Recently, probabilistic methods and statistical learning theory have been shown to provide approximate solutions to "difficult" control problems. Unfortunately, the number of samples required in order to guarantee stringent performance levels may be prohibitively large. This paper introduces bootstrap learning methods and the concept of stopping times to drastically reduce the bound on the number of samples required to achieve a performance level. We then apply these results to obtain more efficient algorithms which probabilistically guarantee stability and robustness levels when designing controllers for uncertain systems.

Keywords

Statistical Learning, Radamacher bootstrap, Robust Control, Sample Complexity, $\mathcal{NP}$-hard problems, Decidability theory.

I. Introduction

It has recently become clear that many control problems are too difficult to admit analytic solutions [12], [15], [17], [56]. New results have also emerged to show that the computational complexity of some "solved" control problems is prohibitive [16], [24], [61]. Many of these (linear and nonlinear) control problems can be reduced to decidability problems or to optimization questions [10], both of which can then be reduced to the question of finding a real vector satisfying a set of (polynomial) inequalities. Even though such questions may be too difficult to answer analytically, or may not be answered exactly given a reasonable amount of computational resources, researchers have shown that we can "approximately" answer these questions "most of the time", and have "high confidence" in the correctness of the answers. In order to fix ideas, we establish the following categorization:

- Which control questions can be answered? Which problems are solvable? This is the realm of Decision Theory, and will give a yes/no answer.
- Which control problems are solvable but difficult? Which problems are solvable but at a prohibitive cost? This is the realm of Computational Complexity Theory, and will tell us which decidable problems are not "practically" solvable.
- What do we do about "approximately" solving (with guaranteed confidence) those problems which are costly to solve exactly? This is the realm of Stochastic Algorithms and Statistical Learning Theory.

Our paper is mainly concerned with the last item above. Many authors have recently advanced the notion of probabilistic methods in control analysis and design. These methods build on the standard Monte Carlo approach (with justifications based on Chernoff Bounds, Hoeffding Inequality, and other elementary probabilistic tools [21], [34], [66]) with ideas advanced during the 1960s and 1970s [63] on the theory of empirical processes and statistical learning. In control theory, some of the original (Monte Carlo) ideas have already been used by Lee and Poolla [45], Ray and Stengel [52], Tempo et al. [6], [59], [60], Barmish et al. [7], [8], [9], [10], Chen and Zhou [18], [19], [20] and by Khargonakar and Tikku [40], to solve *robust analysis* problems while Vidyasagar used learning theory to solve *robust design* problems [66], [68].

Unfortunately, and as acknowledged by the various authors, probabilistic methods, while more efficient than gridding techniques (which suffer from the curse of dimensionality), still require a large number of samples in order to guarantee accurate designs. As an example, Vidyasagar in [68] calculates that more than 2 million samples are needed in order to probabilistically guarantee a certain performance level in a robust control design problem. On the other hand, it was conjectured and verified experimentally that much smaller bounds on the number of samples may be sufficient (tens of thousands instead of millions) to guarantee a certain level of performance [68]. In fact, Vidyasagar in [68] uses 200 samples instead of the millions implied by his bounds, while acknowledging that the theoretical guarantees of accuracy and confidence no longer hold. The question then becomes: what (if any) guarantees are obtained by the smaller number of samples, or more appropriately, is there a smaller bound on the number of samples which can still guarantee the desired level of accuracy and confidence?

This paper answers the last question affirmatively, and does so by invoking different versions of *bootstrap sequential learning* algorithms. For these algorithms, the necessary number of samples (known as the sample complexity of learning) is a random variable whose value is not known in advance and is to be determined in the process of learning. This value is bounded below by the sample size at which the algorithm starts to work, and bounded above by conservative upper bounds of the sample complexity, which are of the same order as the bounds well known in statistical learning theory, used, for instance, by Vidyasagar [66]. This will also lead to the notion of *efficient learning times* which is then used to present our results in a computationally attractive manner.

As will become clearer in the paper, while the decision problem and the optimization problem are related, the bounds are most efficient when one is interested in finding absolute minima, or in solving decision problems, as opposed to finding the near-minima proposed by Vidyasagar [67]. This fact led to some confusion as it turns out that the bounds based on standard Chernoff arguments were actually more efficient (but only for the relatively large values of level parameter α used in [67]) than those rooted in Learning Theory contrary to what was indicated in [67]. We will however illustrate the point that the learning theory bounds are actually required in order to solve the minimization problem if one is interested in finding absolute minima as opposed to level-minima presented later in the paper [67], or when high-dimensional problems are addressed.

The mathematical justification of the methods of learning suggested in this paper relies heavily upon the methods of the empirical processes theory. This theory started in the seminal papers of Vapnik and Chervonenkis [65] and Dudley [28]. The exposition of more recent results on empirical processes can be found in [29] and [62], which also contain a number of deep applications of empirical processes in statistics. The applications of empirical processes to statistical learning problems are discussed in great detail in [25], [63], [64], [66]. The major technical tools used in our paper are concentration inequalities for empirical and related processes. We are using in the current version of the results a relatively old form of these inequalities based on the extension of the classical Hoeffding type bounds to the martingale differences. This extension is due, apparently, to Azuma [5] and it was used very successfully by Yurinskii [69] in the problems of Probability in Banach Spaces. Since then, it has been used in many other applications, including functional limit theorems and empirical processes [41], [42], local theory of Banach spaces [47], combinatorial problems on graphs [46], NP-complete problems [53], and pattern recognition problems [25]. It is clear, however, that further investigation of the properties of sequential learning algorithms would require more advanced and deep versions of concentration inequalities developed in the recent years by Talagrand, see e.g. [57].

The remaining of this paper is divided as follows: section II contains a discussion of generic robust control problems, their difficulty, and their computational complexity. Section III presents an overview of statistical learning methods and section IV contains the bootstrap learning method and its applications to control problems. Section V contains a numerical example illustrating our approach and contrasting it with earlier results, while section VI contains conclusions and an outline for future research. Finally, Appendix A contains the proofs of the main results.

II. Robust Control, Decision Theory, and Computational Complexity

In studying control problems we are led to the conclusion that some robust control problems are actually undecidable. For example, the simultaneous stabilization problem of more than two plants was shown by Blondel [12] to be rationally undecidable using a general model of computing. More examples of such problems may be found in [13].

A. Decision Theory

Most of the control problems we study here are decidable and may be converted to a decision problem relating to the satisfiability of quantified multivariate polynomial inequalities (MPIs) which are then reduced using Tarski's quantifier elimination (QE) theory [58]. These problems include the fixed-structure control design problem for linear and nonlinear systems which remains one of the most practical and difficult problems [31], [56]. In fact, one can argue that most practical control designs involve fixed-structure (and fixed-order) controllers such as PID, or Lead-Lag compensators (see page 113 of [48], and page 3 of [4]). While this makes the control design problem theoretically intractable, it actually reduces some undecidable problems to decidable ones, and fits nicely within the randomized algorithms framework. As an example, the following problems are all decidable using Tarski's decision theory: robust stabilization problems [3], dead-beat control of discrete-time systems [50], Lyapunov stability of polynomial systems [37], and others [2]. The general control problem for an uncertain single input single output (SISO), linear time invariant (LTI) system stated as a decision problem is as follows,

Problem 1: *Given a real rational function* $G(s, X)$*, where* $X = [x_1 \quad x_2 \quad \cdots \quad x_k]$ *is a* k*-dimensional real vector, does there exist an* l*-dimensional real vector* $Y = [y_1 \quad y_2 \quad \cdots \quad y_l]$*,* $y_i \quad [\underline{y_i} \leq y_i \leq \overline{y_i}], \quad 1 \leq i \leq l$*, in the real rational* $C(s, Y)$ *such that for all* $(x_i \quad [\underline{x_i} \leq x_i \leq \overline{x_i}], \quad 1 \leq i \leq k$*, the closed-loop system* $T(s, X, Y)$ *satisfies some performance objectives placed on a scalar performance index* $\Psi(X, Y)$*?*

This is the performance verification problem [10], and it includes the guaranteed-cost design problem. Note that if either X or Y are known, then the problem simplifies to a robust analysis problem. Typical examples are the linear quadratic regulator (LQR), and specialized guaranteed-

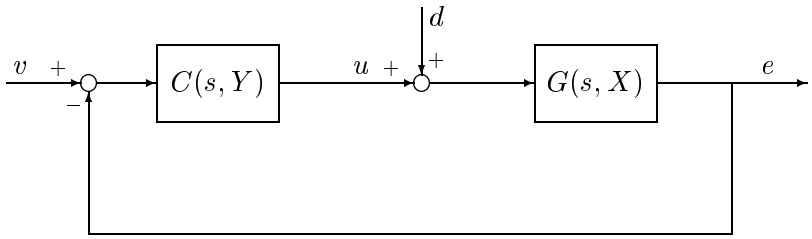

Fig. 1. Feedback Structure for Problem 1

cost problems. In the case where the problem and the performance objective are convex, Linear Matrix Inequalities (LMIs) may be used and the decision control problem is easy. The general decision control problem is however very hard because it leads to a nonlinear, partial differential Hamilton-Jacobi-Bellman (HJB) equation (for non-quadratic performance objectives) which in general is difficult to solve. Researchers in Control Theory have used QE in solving Problem 1 since the 1970's, but the tedious operations made the technique very limited [3]. Later, Collins [22] introduced a theoretically more efficient QE algorithm that uses a cylindrical algebraic decomposition (CAD) approach. However, this algorithm was not capable of effectively handling nontrivial problems. Then, Collins and Hong [23], and Hong [35], [36] introduced a significantly more efficient partial CAD QE algorithm, implemented in the software package QEPCAD. Recently, people have used the QEPCAD software to solve academic, but nontrivial problems [1], [2], [26], [37], [50].

Many robust practical control design problems, for both linear and nonlinear systems, can be reduced to the study of *Boolean* formulae of the type

$$\forall(X \in \mathcal{X})[p_1(X,Y) > 0 \ \wedge \ p_2(X,Y) > 0 \cdots \wedge p_t(X,Y) > 0] \tag{1}$$

where $\forall$ denotes the logic "for all" operator, and $\wedge$ denotes the logic "and" operator. The functions $p_i(X,Y)$ are assumed in this paper to be multivariate polynomial functions, in the components of the vectors X and Y. Note however that our results apply to more general classes of systems (such as nonlinear polynomials) and Boolean formulae. The unquantified variable Y in the formula (1) typically represents controller design parameters, while the quantified variable X represent uncertain plant parameters, state variables (for nonlinear problems) or frequency variables (for linear problems).

B. The Bad News: Computational Complexity of Decidable Problems

If there exists an algorithm which answers a decision problem, the problem is said to be *decidable*. Until recently, it was felt that decidable problems are practically solved and thus not very interesting. The introduction of computational complexity theory has since changed this misconception. Computational complexity theory is often used to establish the tractability or intractability of computational problems, and is concerned with the determination of the intrinsic computational difficulty of these problems and not to any particular algorithm used to solve them [32].

One important concept in this theory is that of a *polynomial-time algorithm*. In practice, such an algorithm can be feasibly implemented on a real computer. This is in contrast to an *exponential-time algorithm*, which is only feasible if the problem being solved is extremely small. Unfortunately, it turns out that QE is at best exponential! [11]

The complexity class $\mathcal{P}$ consists of all decision problems that can be decided in polynomial-time, using a Turing machine model of computation. The simplicity of the Turing machine model appears to make it of little practical value; however, the Church-Turing Thesis holds that the class of problems solvable on a Turing machine in polynomial time is robust across all other reasonable models of computation (including the computers we use).

The complexity class $\mathcal{NP}$ consists of all decision problems that can be decided algorithmically in *nondeterministic* polynomial-time. An algorithm is nondeterministic if it is able to choose or guess a sequence of choices that will lead to a solution, without having to systematically explore all possibilities. This model of computation is not realizable, but it is of theoretical importance since it is strongly believed that $\mathcal{P} \neq \mathcal{NP}$. In other words, these two complexity classes form an important boundary between the tractable (or easy) and intractable (or difficult) problems. A problem is said to be $\mathcal{NP}$-hard if it is as hard as any problem in $\mathcal{NP}$. Thus, if $\mathcal{P} \neq \mathcal{NP}$, the $\mathcal{NP}$-hard problems can only admit deterministic solutions that take an unreasonable (i.e. exponential) amount of time, and they require (unattainable) nondeterminism in order to achieve reasonable (i.e. polynomial) running times.

The central idea used to demonstrate $\mathcal{NP}$-hardness evolves around the $\mathcal{NP}$-complete problems. A problem is said to be $\mathcal{NP}$-complete if every decision problem in $\mathcal{NP}$ is polynomial-time reducible to it. This means that the $\mathcal{NP}$-complete problems are as hard as any decision problem in $\mathcal{NP}$.

Given two decision problems P_1 and P_2, P_1 is said to be polynomial-time reducible to P_2 (written as $\mathrm{P}_1 \leq_p \mathrm{P}_2$), if there exists a polynomial time algorithm R which transforms every input x for P_1 into an equivalent input $R(x)$ for P_2. By equivalent we mean that the answer produced by P_2 on input $R(x)$ is always the same as the answer P_1 produces on input x. Thus, any algorithm which solves P_2 in polynomial time can be used to solve P_1 on input x in polynomial time by simply computing $R(x)$, and then running P_2. In order to show that a particular (control) decision problem P_2 is $\mathcal{NP}$-complete, one starts with a problem P_1 in $\mathcal{NP}$-complete, and attempts to show that $\mathrm{P}_1 \leq_p \mathrm{P}_2$. This shows that P_2 is $\mathcal{NP}$-hard. To complete the proof that P_2 is $\mathcal{NP}$-complete, it must be demonstrated that a candidate solution can be verified in polynomial time. In control theory, researchers have followed this "reduction" method to study the computational difficulty of some decidable problems and many decidable control problems have been shown to be $\mathcal{NP}$-complete (or $\mathcal{NP}$-hard) [16], [24], [49], [51], [61]. A recent overview of the computational complexity of many control problems may be found in [16].

The problem of simultaneous stabilization of N given linear systems with a LTI dynamic compensator is as previously mentioned rationally undecidable for $N > 2$ [12]. However, restricting the stabilizing compensator to be static (or dynamic but of a given order) makes the problem decidable (although inefficiently) using the Tarski approach as discussed before. So the question becomes: how do we deal with decidable but inefficient control problems? And moreover, can we deal with undecidable control problems? We actually have two possibilities in attempting to answer both questions:

1. Limit the class of systems (such as to linear, minimum-phase, passive systems, etc.). This is typically the approach taken by control designers.
2. Soften the goal for the class of systems we are interested in. This is a more recent idea in control pioneered in [18], [52], [55], [66], [70]. An example of goal softening is the randomized algorithms approach discussed next.

A re-formulation of LTI control problems may then be as follows [67],

Problem 2: *Given a closed-loop system $T(s, X, Y)$ with a performance measure $\Psi(X, Y)$, where X, Y are random real-valued vectors, find a vector Y_0, if one exists, of controller parameters which has a high probability of minimizing the expected value with respect to X of an appropriate function $f(X, Y)$ of $\Psi(X, Y)$.*

The related decision problem is to ascertain the existence of a vector Y_0 such that a certain level γ is achieved by $\mathbb{E}f(X,Y)$. Note that our problem has been changed from a deterministic decision problem to a probabilistic optimization problem. Also note that the randomness of X and Y is used to open the door for Monte-Carlo and statistical learning methods. Finally, we have converted a worst-case scenario (guaranteed-cost) into an average-case problem.

In the context of stabilization, let $\Psi(X,Y) = 0$ if $T(s,X,Y)$ is stable and $\Psi(X,Y) = 1$ otherwise. By minimizing $\mathbb{E}f(X,Y)$ we are actually maximizing the volume (or number in case of finite number of plants) which may be stabilized with $C(s,Y_0)$. In fact, let

$$f_Y(X) = f(X,Y) = \begin{cases} 1 & \Psi(X,Y) = 1 \\ 0 & \Psi(X,Y) = 0 \end{cases}$$

and $\mathcal{F} = \{f_Y(\cdot) : Y \in \mathcal{Y}\}$. The purpose of control is to choose Y_0, and thus the corresponding controller $C(s,Y_0)$ to stabilize the maximum number of plants. Note that if the structure and the order of $C(s,Y)$ are fixed, then the problem reduces to finding the set of parameters Y. This objective may be achieved by minimizing the expected value $\mathbb{E}[f_Y(X)]$. An interpretation of the minimization of the expectation $\mathbb{E}[f_Y(X)]$ is that we can then ascertain with confidence $1 - \mathbb{E}[f_{Y_0}(X)]$ that the controller $C(s,Y_0)$ stabilizes a random plant $G(s,X)$.

One limitation of this approach is that in practice, we do not have the necessary information to calculate $\mathbb{E}[f_Y]$ since all we have are sample plants and compensators. Moreover, how do we minimize $\mathbb{E}[f_Y]$ when all we have are the values of f at sample points? In [67], the empirical mean of $f_Y(X)$ is used instead of $\mathbb{E}[f_Y]$ for a given $Y \in \mathcal{Y}$,

$$\frac{1}{n}\sum_{j=1}^{n} f_Y(X_j), \tag{2}$$

which then leaves us with two questions:

1. Will $\frac{1}{n}\sum_{j=1}^{n} f_Y(X_j)$ be a good approximation of $\mathbb{E}[f_Y]$ uniformly in Y as n increases?
2. Will the minimum of $\frac{1}{n}\sum_{j=1}^{n} f_Y(X_j)$, obtained empirically as

$$\min_{1\le i\le m}\left[\frac{1}{n}\sum_{j=1}^{n} f_{Y_i}(X_j)\right]$$

be close to the actual minimum of $\frac{1}{n}\sum_{j=1}^{n} f_Y(X_j)$ as m increases?

It turns out that the first question has been studied thoroughly in the theory of empirical process and statistical learning theory. Minimization of a function defined by equation (2) in particular

is a case of empirical risk minimization as discussed in the next section. Note that there are actually two separate questions to answer: a question of empirical averaging, and a question of empirical minimization. The empirical average question depends on the number n of plants, while the minimization question depends on both the number of plants n and the number of controllers m. Our main results in this paper offer a significant reduction in n but not in m. Our future papers will address the minimization problem and how to reduce m further.

We will next review relevant results from Statistical Learning Theory and randomized algorithms.

III. Overview of Randomized Algorithms and Statistical learning Theory

The basic notions of Probability Theory used in the paper can be found in any textbook on Advanced Probability, see, for instance, [30]. More special results on empirical processes and statistical learning theory can be found in [25], [29], [62], [64], [66]. We present now an overview of standard learning theory concepts and results obtained in [67] along with their application to control problems.

Let $(S, \mathcal{A})$ be a measurable space and let $\{X_n\}_{n\geq 1}$ be a sequence of independent identically distributed (i.i.d) observations in this space with common distribution P. We assume that this sequence is defined on a probability space $(\Omega, \Sigma, \mathbb{P})$. Denote by $\mathcal{P}(S) := \mathcal{P}(S, \mathcal{A})$ the set of all probability measures on $(S, \mathcal{A})$. Suppose $\mathcal{P} \subset \mathcal{P}(S)$ is a class of probability distributions such that $P \in \mathcal{P}$. In particular, if one has no prior knowledge about P, then $\mathcal{P} = \mathcal{P}(S)$. In this case, we are in the setting of *distribution free learning.* One of the central problems of statistical learning theory is *the risk minimization problem.* It is crucial in all cases of learning (standard concept or function learning, regression problems, pattern recognition, etc.). It also plays an important role in randomized (Monte Carlo) algorithms for robust control problems, as has been shown by Vidyasagar [68] and as we will see in this paper. Given a class $\mathcal{F}$ of $\mathcal{A}$-measurable functions f from S into $[0, 1]$ (e.g., decision rules in a pattern recognition problem or performance indices in control problems), the risk functional is defined as

$$R_P(f) := P(f) := \int_S f dP := \mathbb{E}f(X), \ f \in \mathcal{F}.$$

The goal is to find a function f_P that minimizes R_P on $\mathcal{F}$. Typically, the distribution P is unknown (or, as it occurs in many control problems, the integral of f with respect to P is too hard to compute) and the solution of the risk minimization problem is to be based on a sample

$(X_1, \ldots, X_n)$ of independent observations from P. In this case, the goal of statistical learning is more modest: given $\varepsilon > 0, \delta \in (0,1)$, find an estimate $\hat{f}_n \in \mathcal{F}$ of f_P, based on the data $(X_1, \ldots, X_n)$, such that

$$\sup_{P \in \mathcal{P}} \mathbb{P}\{R_P(\hat{f}_n) \geq \inf_{f \in \mathcal{F}} R_P(f) + \varepsilon\} \leq \delta. \tag{3}$$

In other words, one can write that with probability $1-\delta$, $R_P(\hat{f}_n)$ is within ε of $\inf_{f \in \mathcal{F}} R_P(f) = R^*$. Denote by $\tilde{N}^L_{\mathcal{F},\mathcal{P}}(\varepsilon;\delta)$ the minimal number $n \geq 1$ such that for some estimate $\hat{f}_n$ the bound (3) holds, and let $\tilde{N}^U_{\mathcal{F},\mathcal{P}}(\varepsilon;\delta)$ be the minimal number $N \geq 1$ such that for some sequence of estimates $\{\hat{f}_n\}$ and for all $n \geq N$ the bound (3) holds. Let us call the quantity $\tilde{N}^L_{\mathcal{F},\mathcal{P}}(\varepsilon;\delta)$ *the lower sample complexity* and the quantity $\tilde{N}^U_{\mathcal{F},\mathcal{P}}(\varepsilon;\delta)$ *the upper sample complexity* of learning. These quantities show how much data we need in order to guarantee certain accuracy ε of learning with certain confidence level $1-\delta$. Clearly, $\tilde{N}^L_{\mathcal{F},\mathcal{P}}(\varepsilon;\delta) \leq \tilde{N}^U_{\mathcal{F},\mathcal{P}}(\varepsilon;\delta)$, and it is easy to show that the inequality can be strict. The upper sample complexity is used rather frequently in statistical learning theory and is usually referred to simply as the sample complexity. But in this paper we will deal more with the lower sample complexity.

A method of *empirical risk minimization* is widely used in learning theory. Namely, the unknown distribution P is replaced by *the empirical measure* P_n, defined as

$$P_n(A) := \frac{1}{n} \sum_{k=1}^{n} I_A(X_k), \ A \in \mathcal{A}$$

where $I_A(x) = 1$ for $x \in A$ and $I_A(x) = 0$ for $x \notin A$. The risk functional R_P is replaced by the empirical risk R_{P_n}, defined by

$$R_{P_n}(f) := P_n(f) := \int_S f dP_n := \frac{1}{n} \sum_{k=1}^{n} f(X_k), \ f \in \mathcal{F}.$$

The problem is now to minimize the empirical risk R_{P_n} on $\mathcal{F}$, and we let $f_{P_n} \in \mathcal{F}$ be a function that minimizes R_{P_n} on $\mathcal{F}$.

Remark 1: Of course, in general, the minimum in question does not have to exist. It could be replaced in what follows by a random function f_n such that $R_{P_n}(f_n)$ is close enough to $\inf_{f \in \mathcal{F}} R_{P_n}(f)$. For the sake of simplicity, though, we assume throughout the paper that the minimum of R_{P_n} on $\mathcal{F}$ is attained at a random function f_{P_n}, which is properly measurable. We also place proper measurability assumptions on the class $\mathcal{F}$ commonly used in the theory of empirical processes (see, e.g., Dudley [29] or van der Vaart and Wellner [62]).

$\Diamond\Diamond\Diamond$

In what follows, f_{P_n} is used as our learning algorithm, i.e. $\hat{f}_n := f_{P_n}$. Determining the sample complexity of the empirical risk minimization method is definitely one of the central and most challenging problems of statistical learning theory (see, e.g., [25], or Vidyasagar [67] for the relevant discussion in the context of robust control problems). A reasonable upper bound for the sample complexity can be obtained by finding the minimal value of n for which the expected value $\mathbb{E}f(X)$ is approximated uniformly over the class $\mathcal{F}$ by the empirical means with given accuracy ε and confidence level $1-\delta$. More precisely, denote

$$N(\varepsilon,\delta) := N^L_{\mathcal{F},\mathcal{P}}(\varepsilon,\delta) := \min\Big\{n \geq 1 : \sup_{P\in\mathcal{P}} \mathbb{P}\{\|P_n - P\|_{\mathcal{F}} \geq \varepsilon\} \leq \delta\Big\},$$

where $\|\cdot\|_{\mathcal{F}}$ is the sup-norm in the space $\ell^\infty(\mathcal{F})$ of all uniformly bounded functions on $\mathcal{F}$. Let us call the quantity $N(\varepsilon;\delta)$ *the (lower) sample complexity of empirical approximation* on the class $\mathcal{F}$. Then, clearly, $N^L_{\mathcal{F},\mathcal{P}}(\varepsilon/2;\delta) \geq \tilde{N}^L_{\mathcal{F},\mathcal{P}}(\varepsilon;\delta)$. To see this, it is enough to consider the following,

$$0 \leq R_P(f_{P_n}) - \inf_{f\in\mathcal{F}} R_P(f) \leq P(f_{P_n}) - P_n(f_{P_n}) + \inf_{f\in\mathcal{F}} P_n(f) - \inf_{f\in\mathcal{F}} P(f) \leq 2\|P_n - P\|_{\mathcal{F}}. \qquad (4)$$

Unfortunately, the quantity $N^L_{\mathcal{F},\mathcal{P}}(\varepsilon,\delta)$ is itself unknown for most of the nontrivial examples of function classes, and only rather conservative upper bounds for this quantity are available. These bounds are expressed in terms of various entropy characteristics and combinatorial quantities, such as VC-dimensions, which themselves are not always known precisely and are replaced by their upper bounds [67].

Going back to our control motivation, we note that our problem involves also the finding of the minimum of a certain performance objective or more precisely, finding the controller parameters which correspond to such minimum. This is the second separate question mentioned at the end of Section II and refers to the optimization part of the problem which we approach in the same manner as Vidyasagar.

In [67], Vidyasagar introduced the following types of minima, in order to use statistical learning theory to design fixed-order robust controllers, which minimize the performance index in Problem 2.

Definition 1: Let $R : \mathcal{Y} \longrightarrow \mathbb{R}$ and $\varepsilon > 0$ be given. A number $R_0 \in \mathbb{R}$ is said to be an *approximate near minimum of R to accuracy ε* if

$$\Big|R_0 - \inf_{Y\in\mathcal{Y}} R(Y)\Big| \leq \varepsilon$$

■

Definition 2: Suppose $R : \mathcal{Y} \longrightarrow \mathbb{R}$, Q is a given probability measure on $\mathcal{Y}$, and $\alpha > 0$ be given. A number $R_0 \in \mathbb{R}$ is a *probable near minimum of R to level α* if there exists a measurable set $\mathcal{S} \subseteq \mathcal{Y}$ with $Q(\mathcal{S}) \leq \alpha$ such that

$$\inf_{Y \in \mathcal{Y}} R(Y) \leq R_0 \leq \inf_{Y \in \mathcal{Y} \setminus \mathcal{S}} R(Y).$$

where $\mathcal{Y} \setminus \mathcal{S}$ is the complement of the set $\mathcal{S}$ in $\mathcal{Y}$.

■

Definition 3: Suppose $R : \mathcal{Y} \longrightarrow \mathbb{R}$, Q is a given probability measure on $\mathcal{Y}$, and $\alpha > 0$, $\varepsilon > 0$ be given. A number $R_0 \in \mathbb{R}$ is a *probably approximate near minimum of R to accuracy ε and level α* if there exists a measurable set $\mathcal{S} \subseteq \mathcal{Y}$ with $Q(\mathcal{S}) \leq \alpha$ such that

$$\inf_{Y \in \mathcal{Y}} R(Y) - \varepsilon \leq R_0 \leq \inf_{Y \in \mathcal{Y} \setminus \mathcal{S}} R(Y) + \varepsilon.$$

■

Note in particular that the last 2 types of minima while useful in practice may not give an accurate picture of the performance of the closed-loop control system. It is the relaxation to the level α that makes Algorithm 1 below at first glance much more efficient than later ones. Finally, let us define a version of probably approximate near minima in the case of a stochastic process R (say, $R := R_{P_n}$, see the definition above) as follows.

Definition 4: Suppose that $R : \mathcal{Y} \to \mathbb{R}$ is a stochastic process, that Q is a given probability measure on $\mathcal{Y}$, and that $\alpha \in (0, 1)$, $\delta \in (0, 1)$ and $\varepsilon > 0$ are given. A number R_0 is a *probably approximate near minimum of R with confidence $1 - \delta$, level α and accuracy ε*, if

$$\mathbb{P}\left\{ \inf_{Y \in \mathcal{Y}} R(Y) - \varepsilon \leq R_0 \leq \inf_{Y \in \mathcal{Y} \setminus \mathcal{S}} R(Y) + \varepsilon \right\} \geq 1 - \delta$$

with some measurable set $\mathcal{S} \subseteq \mathcal{Y}$ such that $Q(\mathcal{S}) \leq \alpha$.

■

An interpretation of definitions 2, 3, 4 is that we are not searching for the minimum over all of the set $\mathcal{Y}$ but only over its subset $\mathcal{Y} \setminus \mathcal{S}$, where $\mathcal{S}$ has a small measure (at most α). Unless the actual

infimum R^* is attained in the exceptional set $\mathcal{S}$, R_0 is within ε from the actual infimum with confidence $1-\delta$. It is exactly this goal softening that gets around the computational difficulty of these problems [70]. Although using Monte Carlo type minimization, it is unlikely to obtain a better estimate of R^* than R_0 (since the chances of getting into the set $\mathcal{S}$ are small), nothing can be said in practice about the size of the difference $R_0 - R^*$. The following two stochastic algorithms were then presented in [67] to solve the problem of designing robust controllers.

Algorithm 1: Given:

- Spaces $\mathcal{X}$ and $\mathcal{Y}$,
- Probability measures P on $\mathcal{X}$ and Q on $\mathcal{Y}$,
- A measurable function $f : \mathcal{X} \times \mathcal{Y} \longrightarrow [0,1]$, and
- An accuracy parameter $\varepsilon \in (0,1)$, a level parameter $\alpha \in (0,1)$, and a confidence parameter $\delta \in (0,1)$.

Let

$$R_P(\cdot) = \mathbb{E}_P[f(X,\cdot)]$$

and let

$$\begin{aligned} m &\geq \frac{\log(2/\delta)}{\log[1/(1-\alpha)]} \\ n &\geq \frac{1}{2\varepsilon^2}\log\frac{4m}{\delta}. \end{aligned}$$

Generate i.i.d. samples $Y_1, Y_2, \cdots, Y_m \in \mathcal{Y}$ from the distribution Q and $X_1, X_2, \cdots, X_n \in \mathcal{X}$ from the distribution P. Then let,

$$\begin{aligned} R_{P_n}(\cdot) &= \frac{1}{n}\sum_{j=1}^{n} f(X_j,\cdot); \\ R_0 &= \min_{1\leq i\leq m} R_{P_n}(Y_i). \end{aligned}$$

Then with confidence at least $1-\delta$, R_0 is a probably approximate near minimum of $R_P(Y)$ to level α and accuracy ε.

Note that algorithm is very general and apply to any measurable function f. Also, note that the probability measures P and Q may be given (such as the case of Normal distribution for

the system's parameters) or just chosen (as uniform) since the results will are independent of a particular probability distribution. Finally, note that the bound on n depends on m and thus implicitly on the level α. Vidyasagar in [67] then proposes a more "efficient" algorithm as follows.

Algorithm 2: Given:

- Sets $\mathcal{X}$ and $\mathcal{Y}$,
- Probability measures P on $\mathcal{X}$ and Q on $\mathcal{Y}$,
- A measurable function $f : \mathcal{X} \times \mathcal{Y} \longrightarrow [0, 1]$, and
- An accuracy parameter $\varepsilon \in (0, 1)$, a level parameter $\alpha \in (0, 1)$, and a confidence parameter $\delta \in (0, 1)$.

Let $R_P(\cdot) = \mathbb{E}_P[f(X, \cdot)]$ and denote $\mathcal{F} := \{f_Y : Y \in \mathcal{Y}\}$,

$$\begin{aligned} R_{P_n}(\cdot) &= \frac{1}{n}\sum_{j=1}^{n} f(X_j, \cdot) \\ q(n, \varepsilon, \mathcal{F}) &= \mathbb{P}\{\sup_{Y \in \mathcal{Y}} |R_{P_n}(Y) - R_P(Y)| > \varepsilon\}. \end{aligned}$$

Then, choose n and m such that

$$\begin{aligned} m &\geq \frac{\log(2/\delta)}{\log[1/(1-\alpha)]} \\ q(n, \varepsilon, \mathcal{F}) &\leq \delta/2 \end{aligned}$$

and generate i.i.d. samples $Y_1, Y_2, \cdots, Y_m \in \mathcal{Y}$ from the distribution Q and $X_1, X_2, \cdots, X_n \in \mathcal{X}$ from the distribution P. Then let,

$$R_0 = \min_{1 \leq i \leq m} R_{P_n}(Y_i)$$

Then with confidence at least $1 - \delta$, R_0 is a probably approximate near minimum of $R_P(Y)$ to level α and accuracy ε.

To guarantee the existence of n such that $q(n, \varepsilon, \mathcal{F}) \leq \delta/2$, in Algorithm 2 one can assume that $\mathcal{F}$ is a Glivenko-Cantelli class for P (see [29], [62] for the definition). The UCEM property considered in [66] means that for all $\varepsilon > 0$ $q(n, \varepsilon, \mathcal{F}) \to 0$ as $n \to \infty$; it is equivalent to the Glivenko-Cantelli property of the class $\mathcal{F}$. Note that in Algorithm 2, the bound on the quantity $q(n, \varepsilon, \mathcal{F})$ is no longer dependent on m. This along with other considerations have led Vidyasagar

[67] to present Algorithm 2 as more efficient than Algorithm 1. It turns out that the reverse is actually true. Namely, for the values of α that are not particularly small and that were used in [68], Algorithm 1 which only relies on standard Chernoff bounds [21] is much more computationally efficient than Algorithm 2 which, as described in [66], requires the introduction of modern tools of statistical learning theory. We note however, that in the multidimensional situation, the simple Monte Carlo scheme of minimization used in Algorithm 2 can be very misleading and the empirical minimum can be much larger than the true minimum with probability practically equal to 1. Imagine, for instance, that the function to be minimized is defined on the unit ball $B_l := \{x : |x| \leq 1\}$ in $\mathbb{R}^l$, say, with $l = 100$, and is given by the following expression:

$$R(Y) := \sqrt{y_1^2 + y_2^2 + \ldots + y_l^2},\ Y = (y_1, \ldots, y_l).$$

Suppose also that the distribution Q on B_l is uniform. Let $r \in (0,1)$ and let $(Y_1, \ldots, Y_m)$ be an i.i.d. sample from Q. Then the probability that all the points in the sample are outside the ball rB_l is equal to $(1 - r^l)^m$. This implies that

$$\mathbb{P}\{\min_{1 \leq i \leq m} R(Y_i) \geq r\} \geq (1 - r^l)^m.$$

For instance, if one takes $r = 1/4$ and $m = 2^{100}/10^{10} > 10^{20}$, one would get that

$$\mathbb{P}\{\min_{1 \leq i \leq m} R(Y_i) \geq 1/4\} \geq 1 - 10^{-10}.$$

So, even with such enormously large sample sizes the Monte Carlo search for minimum would not come even close to the true minimum of R, which is equal to 0! With more reasonable sample sizes, the empirical minimum will be very close to 1 (which, in fact, is the maximum of R on B_l) with probability close to 1. It looks like almost any other method of minimization would do better in such an example than Monte Carlo does. In order to make the method work, one has to choose α extremely small, and in this case, of course, the computational efficiency of the Algorithm 1 disappears. This problem is of course known in problems of high-dimensional geometry [33] and will manifest itself as the "fragility" of high-dimensional controllers [39]. In particular, if one is required to use a high-dimensional controller, as usually happens when attempting to meet stringent H_∞ objectives, then it is conceivable that a very small α is required in order to find an appropriate coefficient vector Y. Therefore, in such cases, the number of sample controllers m will increase as one is searching for a probably approximate near minimum outside the set of small measure α, and resulting in performance measures which can grow unacceptably large.

In such situations, more efficient methods of minimization should be used and their justification would heavily rely on statistical learning theory (which allows us to determine the sample complexity of risk minimization regardless of the particular minimization algorithm). The Monte Carlo minimization scheme suggested in [66], [67], and used in Algorithms 1, and 2, can definitely be used in preliminary studies of new learning algorithms, and it could in many cases provide satisfactory results in control design. In fact, we use this approach in the next section in combination with our new sequential learning algorithms.

Sufficient conditions for satisfying Glivenko-Cantelli (UCEM) property, which are convenient for the purposes of control theory, can be formulated in terms of the finiteness of VC-dimensions or P-dimensions of the class $\mathcal{F}$, [66], [67].

Definition 5: Let $\mathcal{C}$ be a family of subsets of $\mathcal{X}$. A finite set $F = \{x_1, \cdots x_n\} \subset \mathcal{X}$ is shattered by $\mathcal{C}$, if for every subset B of the 2^n subsets of F, there exists a set $A \in \mathcal{C}$ such that $F \cap A = B$. The Vapnik-Chervonenkis dimension of $\mathcal{C}$ denoted $VC\text{-}dim(\mathcal{C})$ is the largest integer n such that there exists a set F of cardinality n shattered by $\mathcal{C}$.

■

Given a class $\mathcal{F}$ of functions mapping $\mathcal{X}$ into $\{0,1\}$, one can consider the class of sets $\mathcal{C} := \{\{x : f(x) = 1\} : f \in \mathcal{F}\}$ and define the VC-dimension of $\mathcal{F}$ as $VC\text{-}dim(\mathcal{C})$. It will be also denoted $VC\text{-}dim(\mathcal{F})$. The role of P-dimension (see e.g. [66]) is similar in the case of more general classes of functions. In particular, one can consider the class $\mathcal{F}_{k,l,r,t}$ arising from our MPIs and defined as follows. Given polynomials $p_1(X,Y), \ldots, p_t(X,Y)$ on $\mathbb{R}^k \times \mathbb{R}^l$ of degree $\leq r$ (with respect to Y), consider all the Boolean formulae obtained from expressions "$p_j(X,Y) > 0$", $j = 1, \ldots, t$ using the standard logical operations $\vee, \wedge, \neg$. Let $\Phi_{k,l,r,t}$ be the set of all such formulae. Each formula $\phi \in \Phi_{k,l,r,t}$ defines the function $f := f_\phi$ that takes value 1 if the formula is true and value 0 otherwise. We set $\mathcal{F}_{k,l,r,t} := \{f_\phi : \phi \in \Phi_{k,l,r,t}\}$. This class can be used to describe the control decidability questions.

We then have the following theorems that go back to the original work of Vapnik and Chervonenkis [64], [65] and that were used in [67].

Theorem 1: Let $\mathcal{F}$ be a family of measurable functions from $\mathcal{X}$ into $\{0,1\}$ and suppose that

VC-dim$(\mathcal{F}) \leq d < \infty$. Then, $\mathcal{F}$ has the UCEM property and moreover,

$$q(n, \varepsilon, \mathcal{F}) \leq 4 \left(\frac{2en}{d} \right)^d \exp(-n\varepsilon^2/8); \quad \forall n, \varepsilon$$

■

This then leads to the following bound on the sample complexity of empirical approximation on the class $\mathcal{F}$.

Theorem 2: Let $\mathcal{F}$ be a family of measurable functions from $\mathcal{X}$ into $\{0, 1\}$ and suppose that VC-dim$(\mathcal{F}) \leq d < \infty$. Let P be an arbitrary probability measure on $\mathcal{X}$, and let $\varepsilon, \delta \in (0, 1)$ be arbitrary constants. Then, $q(n, \varepsilon, \mathcal{F}) \leq \delta$ if

$$n \geq \max \left\{ \frac{16}{\varepsilon^2} \log \frac{4}{\delta}, \frac{32d}{\varepsilon^2} \log \frac{32e}{\varepsilon^2} \right\}$$

■

The next theorem gives an upper bound for the VC-dimension of the class $\mathcal{F}_{k,l,r,t}$ and is due essentially to Karpinski and Macintyre [38]. We cite it from [66].

Theorem 3: The following upper bound holds:

$$VC - dim(\mathcal{F}_{k,l,r,t}) \leq 2l \log(4ert).$$

■

Following Vidyasagar [67], [68], our initial purpose is to explore the utility of statistical learning theory in the efficient design of robust controllers for linear uncertain systems. Throughout the discussion we will refer generically to a real rational plant $G(s, X)$ and a real rational controller $C(s, Y)$ defined as:

$$\begin{aligned} G(s, X) &= \frac{n_G(s, X)}{d_G(s, X)}, \ X \in \mathcal{X}, \\ C(s, Y) &= \frac{n_C(s, Y)}{d_C(s, Y)}, \ Y \in \mathcal{Y} \end{aligned}$$

where $\mathcal{X} \subseteq \mathbb{R}^k, \mathcal{Y} \subseteq \mathbb{R}^l$, and n_G, d_G, n_C, d_C are polynomial in their arguments and where $\alpha_s \stackrel{\triangle}{=}$ degree of $G(s, X)$ in s, $\beta_s \stackrel{\triangle}{=}$ degree of $C(s, Y)$ in s, and $\beta_Y \stackrel{\triangle}{=}$ Maximum degree of Y in $n_C(s, Y)$ or $d_C(s, Y)$.

The methods to be used require us to sample from a set of possible plants and controllers. Since the actual distributions are unknown, we postulate uniform measures for both $\mathcal{X}$ and $\mathcal{Y}$ (assuming that these sets are bounded).

An estimate of the number of samples is first obtained using the results cited in Algorithms 1 and 2. However, before we can estimate sample size, we must specify values for the following parameters:

$$\begin{aligned} \alpha &\stackrel{\triangle}{=} \text{the level parameter} \\ \delta &\stackrel{\triangle}{=} \text{is related to the final confidence level } (1-\delta) \\ \varepsilon &\stackrel{\triangle}{=} \text{the bound on the final accuracy} \end{aligned}$$

Sample sizes n and m may then be estimated:

$$\begin{aligned} m &\geq \frac{\log(2/\delta)}{\log[1/(1-\alpha)]} \\ n &\geq \frac{1}{2\varepsilon^2}\log\frac{4m}{\delta} \quad \text{for Algorithm 1} \\ n &\geq \max\left\{\frac{16}{\varepsilon^2}\log\frac{4}{\delta}, \frac{32d}{\varepsilon^2}\log\frac{32e}{\varepsilon^2}\right\} \quad \text{for Algorithm 2} \end{aligned}$$

Note here that d is either the VC or the P-dimension of the problem [66].

Example 1: The following plant and controller will be used in the first example, to illustrate the fact that Algorithm 1 is indeed more efficient than Algorithm 2:

$$\begin{aligned} G(s,X) &= \frac{X_1}{1-s/X_2}, \quad 0.8 \leq X_1, X_2 \leq 1.25 \\ K(s,Y) &= Y_1 \end{aligned}$$

We attempt to choose a controller that stabilizes the plant given above. We begin by sampling X_1 and X_2 uniformly between 0.8 and 1.25, and sampling Y_1 between -100 and 100. Thus, $k=2$, $l=1$, $\alpha_s=1$, $\beta_s=0$, $\beta_Y=1$. We further choose $\alpha=0.05$, $\varepsilon=0.1$, and $\delta=0.05$. The required number of samples then becomes according to Algorithm 2:

$$\begin{aligned} m &= 72 \text{ samples of } Y_1 \\ n &= 199,862 \text{ samples of } X_1, X_2 \end{aligned}$$

For comparison, we calculate the estimate from Algorithm 1: $n \geq \frac{1}{2\varepsilon^2}\log\frac{4m}{\delta}$. From this we obtain:

$$n = 433.$$

Note that this result is much smaller than the prior estimate obtained based on Algorithm 2.

$\triangle\triangle\triangle$

Example 2: Consider the linearized model of a CH-47 tandem-rotor helicopter in horizontal motion about a nominal airspeed of 40 knots, as discussed in [27]. The model parameters are given by

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

where

$$A = \begin{bmatrix} -0.02 & 0.005 & 2.4 & -32 \\ -0.14 & 0.44 & -1.3 & -30 \\ 0 & 0.018 & -1.6 & 1.2 \\ 0 & 0 & 1 & 0 \end{bmatrix}$$

$$B = \begin{bmatrix} 0.14 & -0.12 \\ 0.36 & -8.6 \\ 0.35 & 0.009 \\ 0 & 0 \end{bmatrix}; \quad C = \begin{bmatrix} 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 57.3 \end{bmatrix}$$

The incremental outputs are

- y_1 is the vertical velocity (knots/hr)
- y_2 is the pitch altitude (radians)

and the inputs are

- u_1 is the collective rotor thrust
- u_2 is the differential collective rotor thrust

Let the output feedback be $u = Ky$ and let all the coefficients of A, B, C (except for the zero and unity terms) be perturbed randomly over ranges of about half of their magnitudes. The controller gain matrix contains 4 parameters Y_1, Y_2, Y_3, Y_4. Similarly to the previous example, by

choosing $\alpha = 0.05$, $\varepsilon = 0.1$, and $\delta = 0.05$, we generate 72 controllers and 433 plants, and choose the controller gain leading to robust stability and to the minimization of the H_2 performance for the nominal system. The constant gain

$$K = \begin{bmatrix} -12.7177 & -45.0824 \\ 63.5123 & 25.9144 \end{bmatrix}$$

leads to the H_2 performance of 0.339458, and to the closed-loop eigenvalues at -0.02, $-0.88 \pm 28.15j$, -550.18 for the nominal system. This example illustrates that the dimensions of the system and of the controller are not critical to the number of samples, if one uses Algorithm 1, but will influence the number of plants n through the dimension d for Algorithm 2.

$\triangle\triangle\triangle$

IV. The Good News: Sequential Learning Algorithms

In this section, we present sequential algorithms for a general problem of empirical risk minimization. They are designed to overcome some of the difficulties encountered with the standard learning methods of Section III. This approach does not depend on the explicit calculation of the VC-dimension, although its finiteness remains critical to the termination of the design algorithm, in the distribution-free learning case. The sequential algorithms chosen are based on *Rademacher bootstrap* although other bootstrap techniques, developed in statistics (for instance, standard Efron bootstrap or various versions of weighted bootstrap), can also be adopted for our purposes. An important feature of our approach is the randomness of the sample size for which a given accuracy of learning is achieved with a guaranteed probability. Thus, the sample complexity of our method of learning is rather a random variable. Its value is not known in advance and is to be determined in the process of learning. The lower bound for this random variable is the value of the sample size which the sequential learning algorithm starts working with. The upper bounds for the random sample complexity are of the same order of magnitude as the standard conservative upper bounds for the sample complexity of empirical risk minimization algorithms. Thus, *in the worst case,* the sequential method of learning would take as much time (up to a numerical constant) as the standard methods do.

We start with several basic definitions. The proofs of all statements of this section can be found in the Appendix.

Definition 6: Let $\{\Sigma_n\}_{n\geq 1}$ be a filtration of σ-algebras (i.e. for all $n \geq 1$ $\Sigma_n \subset \Sigma_{n+1}$) such that $\Sigma_n \subset \Sigma, n \geq 1$ and X_n is Σ_n-measurable. Less formally, Σ_n consists of the events that occur by time n (in particular, the value of random variable X_n is known by time n). A random variable τ, taking positive integer values, will be called *a stopping time* if and only if (iff), for all $n \geq 1$, we have $\{\tau = n\} \in \Sigma_n$. In other words, the decision whether $\tau \leq n$, or not, depends only on the information available by time n.

■

Given $\varepsilon > 0$ and $\delta \in (0,1)$, let $\bar{n}(\varepsilon, \delta)$ denote the initial sample size of our learning algorithms. We assume that $\bar{n}$ is a non-increasing function in both ε and δ. Denote by $\mathcal{T}(\varepsilon, \delta) := \mathcal{T}_{\mathcal{F},\mathcal{P}}(\varepsilon, \delta)$ the set of all stopping times τ such that $\tau \geq \bar{n}(\varepsilon; \delta)$ and

$$\sup_{P\in\mathcal{P}} \mathbb{P}\{\|P_\tau - P\|_{\mathcal{F}} \geq \varepsilon\} \leq \delta.$$

If now $\tau \in \mathcal{T}(\varepsilon, \delta)$ and $\hat{f} := f_{P_\tau}$ is a function that minimizes the empirical risk based on the sample $(X_1, \ldots, X_\tau)$ then a bound similar to (4) immediately implies that

$$\sup_{P\in\mathcal{P}} \mathbb{P}\Big\{R_P(f_{P_\tau}) \geq \inf_{f\in\mathcal{F}} R_P(f) + 2\varepsilon\Big\} \leq \delta.$$

The questions, though, are how to construct a stopping time from the set $\mathcal{T}(\varepsilon, \delta)$, based only on the available data (without using the knowledge of P) and which of the stopping times from this set is best used in the learning algorithms. The following definition will be useful in this connection.

Definition 7: A parametric family of stopping times $\{\nu(\varepsilon, \delta) : \varepsilon > 0, \delta \in (0,1)\}$ is called *strongly (statistically) efficient* for the class $\mathcal{F}$ with respect to $\mathcal{P}$ iff there exist constants $K_1 \geq 1, K_2 \geq 1$ and $K_3 \geq 1$ such that for all $\varepsilon > 0$ and $\delta \in (0,1)$

$$\nu(\varepsilon, \delta) \in \mathcal{T}(K_1\varepsilon, \delta)$$

and for all $\tau \in \mathcal{T}(\varepsilon, \delta)$

$$\sup_{P\in\mathcal{P}} \mathbb{P}\{\nu(K_2\varepsilon, \delta) > \tau\} \leq K_3\delta.$$

■

Thus, using strongly efficient stopping time $\nu(\varepsilon; \delta)$ allows one to solve the problem of empirical approximation with confidence $1 - \delta$ and accuracy $K_1\varepsilon$. With probability at least $1 - K_3\delta$, the

time required by this algorithm is less than the time needed for *any* sequential algorithm of empirical approximation with accuracy ε/K_2 and confidence $1-\delta$.

Definition 8: We call a family of stopping times $\{\nu(\varepsilon,\delta) : \varepsilon > 0, \delta \in (0,1)\}$ *weakly (statistically) efficient* for the class $\mathcal{F}$ with respect to $\mathcal{P}$ iff there exist constants $K_1 \geq 1, K_2 \geq 1$ and $K_3 \geq 1$ such that for all $\varepsilon > 0$ and $\delta \in (0,1)$

$$\nu(\varepsilon,\delta) \in \mathcal{T}(K_1\varepsilon,\delta)$$

and

$$\sup_{P\in\mathcal{P}} \mathbb{P}\{\nu(K_2\varepsilon,\delta) > N(\varepsilon;\delta)\}\} \leq K_3\delta.$$

■

Using weakly efficient stopping time $\nu(\varepsilon;\delta)$ also allows one to solve the problem of empirical approximation with accuracy $K_1\varepsilon$ and confidence $1-\delta$. With probability at least $1-K_3\delta$, the time required by this algorithm, is less than the sample complexity of empirical approximation with accuracy ε/K_2 and confidence $1-\delta$.

Note that, under the assumption $N(\varepsilon;\delta) \geq \bar{n}(\varepsilon;\delta)$, we have $N(\varepsilon,\delta) \in \mathcal{T}(\varepsilon,\delta)$. Hence, any strongly efficient family of stopping times is also weakly efficient. The converse to this statement is not true.

Proposition 1: There exists a weakly efficient family of stopping times that is not strongly efficient.

Proof: See Appendix -A. ■

We show below how to construct efficient stopping times for empirical risk minimization problems. The construction is based on a version of bootstrap. Let $\{r_n\}_{n\geq 1}$ be *a Rademacher sequence* (i.e. a sequence of i.i.d. random variables taking values $+1$ and -1 with probability $1/2$ each). We assume, in addition, that this sequence is independent of the observations $\{X_n\}_{n\geq 1}$. Suppose that (with $\lfloor\cdot\rfloor$ denoting the floor of the argument)

$$\bar{n}(\varepsilon,\delta) \geq \left\lfloor \frac{4}{\varepsilon^2}\log(\frac{2}{\delta(1-e^{-\varepsilon^2/4})}) \right\rfloor + 1.$$

Let

$$\nu(\varepsilon,\delta) := \nu_{\mathcal{F}}(\varepsilon,\delta) := \min\{n \geq \bar{n}(\varepsilon,\delta) : \|n^{-1}\sum_{j=1}^{n} r_j\delta_{X_j}\|_{\mathcal{F}} \leq \varepsilon\}.$$

where $\delta_x(f) := f(x)$. Note that for all $\varepsilon > 0$ and for all $\delta \in (0,1)$, $\nu(\varepsilon,\delta)$, is a stopping time and it can be computed by Monte Carlo simulation of the sequence $\{r_j\}_{j\geq 1}$. The finiteness with probability 1 of the stopping time $\nu(\varepsilon;\delta)$ (and other stopping times, defined below) can be shown to follow from the Glivenko-Cantelli property for the class $\mathcal{F}$ (also referred to as UCEM property, see section III).

Theorem 4: $\{\nu_{\mathcal{F}}(\varepsilon,\delta) : \varepsilon > 0, \delta \in (0,1)\}$ is a strongly efficient family of stopping times for any class $\mathcal{F}$ of measurable functions from S into $[0,1]$ with respect to the set $\mathcal{P}(S)$ of all probability distributions.

■

The proof of Theorem 4 is based on proposition 2 below. Let

$$\Delta_n := \|P_n - P\|_{\mathcal{F}}.$$

Proposition 2: For all $\varepsilon > 0$ and $\delta > 0$

1.

$$\sup_{P\in\mathcal{P}(S)} \mathbb{P}\{\Delta_{\nu(\varepsilon,\delta)} \geq 5\varepsilon\} \leq \delta$$

and

2.

$$\sup_{P\in\mathcal{P}(S)} \mathbb{P}\{\min_{\bar{n}(\varepsilon,\delta)\leq n<\nu(6\varepsilon,\delta)} \Delta_n < \varepsilon\} \leq \delta.$$

■

The proof of the proposition may be found in Appendix -A. It immediately follows from the Proposition 2 that in Theorem 4, $\{\nu(\varepsilon,\delta) : \varepsilon > 0, \delta \in (0,1)\}$ is strongly efficient with $K_1 = 5, K_2 = 6$ and $K_3 = 2$.

A. Other versions of the sequential algorithm

The initial time of the previous algorithm could be too large if ε is very small. Here we construct another version of sequential risk minimization algorithm with smaller initial time.

Define

$$\nu(\varepsilon,\delta) := \nu_{\mathcal{F}}(\varepsilon,\delta) := \min\{n : \|n^{-1}\sum_{j=1}^{n} r_j\delta_{X_j}\|_{\mathcal{F}} \leq \varepsilon,\ n := n_k := 2^k\bar{n}(\varepsilon,\delta),\ k = 0,1,\ldots\}.$$

Theorem 5: Suppose that

$$\bar{n}(\varepsilon,\delta) \geq \left\lfloor \frac{4}{\varepsilon^2}\log(\frac{4}{\delta}) \right\rfloor + 1.$$

Then, for all $\varepsilon > 0, \delta \in (0,1)$,

1. $\nu(\varepsilon;\delta) \in \mathcal{T}(K_1\varepsilon;\delta)$ with $K_1 = 5$.

2. Moreover, suppose that

$$N(\varepsilon,\delta) \geq \bar{n}(\varepsilon,\delta) \geq \left\lfloor \frac{4}{\varepsilon^2}\log(\frac{4}{\delta}) \right\rfloor + 1.$$

Then $\{\nu_{\mathcal{F}}(\varepsilon,\delta) : \varepsilon > 0, \delta \in (0,1/2)\}$ is a weakly efficient family of stopping times for any class $\mathcal{F}$ of measurable functions from S into $[0,1]$ with respect to the set $\mathcal{P}(S)$ of all probability distributions on S.

Proof: See Appendix -A. ■

The next proposition shows that if the family of stopping times defined above starts too late (namely, after the time $N(\varepsilon;\delta)$), then the stopping time is close to the initial time with high probability.

Proposition 3: Suppose that

$$\bar{n}(\varepsilon,\delta) \geq \left\lfloor \frac{4}{\varepsilon^2}\log(\frac{4}{\delta}) \right\rfloor + 1$$

and

$$12/\varepsilon \leq N(\varepsilon,\delta) \leq \bar{n}(\varepsilon,\delta).$$

Then, there exist constants $K_1 \geq 1, K_2 \geq 1$ such that

$$\sup_{P\in\mathcal{P}(S)} \mathbb{P}\{\nu_{\mathcal{F}}(K_1\varepsilon;\delta) > K_2\bar{n}(\varepsilon;\delta)\} \leq \delta. \tag{5}$$

Proof: See Appendix -A. ■

Based on the randomized algorithms introduced in section III, and on the sequential learning algorithms of this section, a probably approximate near minimum of f with confidence $1-\delta$, level α and accuracy ε, can be found with the following algorithm.

Algorithm 3: Given:

- Sets $\mathcal{X}$ and $\mathcal{Y}$,
- Probability measures P on $\mathcal{X}$ and Q on $\mathcal{Y}$,
- A measurable function $f : \mathcal{X} \times \mathcal{Y} \longrightarrow [0, 1]$, and
- An accuracy parameter $\varepsilon \in (0, 1)$, a level parameter $\alpha \in (0, 1)$, and a confidence parameter $\delta \in (0, 1)$.

Let $R_P(\cdot) = \mathbb{E}_P[f(X, \cdot)]$ and

$$R_{P_n}(\cdot) = \frac{1}{n} \sum_{j=1}^{n} f(X_j, \cdot)$$

Then,

1. Choose m independent controllers with parameters having distribution Q where

$$m \geq \frac{\log(2/\delta)}{\log[1/(1-\alpha)]}$$

2. Choose n independent plants with parameters having distribution P, where

$$n = \left\lfloor \frac{4K_1^2}{\varepsilon^2} \log\left(\frac{8}{\delta}\right) \right\rfloor + 1$$

with $K_1 = 5$

3. Evaluate the stopping variable

$$\gamma = \max_{1 \leq i \leq m} \left| \frac{1}{n} \sum_{j=1}^{n} r_j f(X_j, Y_i) \right|$$

where r_j are *Rademacher* random variables, i.e. independent identically distributed random variables (also independent of the plant sample) taking values $+1$ and -1 with probability $1/2$ each. If $\gamma > \frac{\varepsilon}{K_1}$, add n more independent plants with parameters having distribution P to the plant samples, set $n := 2n$ and repeat step 3

4. Choose the controller which minimizes the cost function R_{P_n}. Then with confidence at least $1 - \delta$, this controller minimizes R_P to a level α and accuracy ε.

Note that Algorithm 3 corresponds to Theorem 5 and other variations on this algorithm are possible.

V. Applications To Control Design

Example 3: In this example we consider the control problem presented by Vidyasagar in [68] and solved via randomized algorithms. This will allow us to illustrate our method and to compare it to the one proposed in [68].

The example concerns the design of an inner-loop controller for the longitudinal axis of an aircraft. The problem is to minimize the weighted sensitivity function over a certain set of uncertain plants, given some constraints on the nominal plant. For further details, the reader is referred to [68].

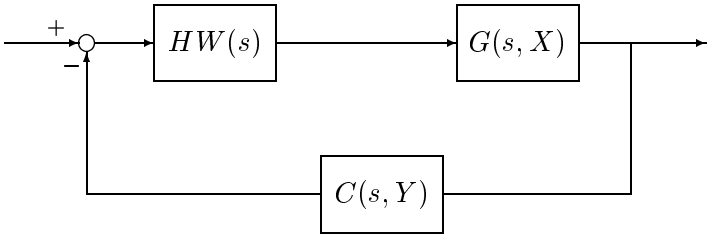

Fig. 2. The closed-loop system

The closed-loop system is shown in Figure 2. The plant $G(s,X)$ is in the form

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

where

$$A = \begin{bmatrix} Z_\alpha & 1 - Z_q \\ M_\alpha & M_q \end{bmatrix},\ B = \begin{bmatrix} Z_{\delta e} \\ M_{\delta e} \end{bmatrix},\ C = \begin{bmatrix} 1 & 0 \\ 0 & 1 \end{bmatrix}$$

The parameters of the matrices have Gaussian distribution with means and standard deviations as in Table I. In the following, we let $X = [Z_\alpha \ \ Z_q \ \ M_\alpha \ \ M_q \ \ Z_{\delta e} \ \ M_{\delta e}]^T$.

The transfer function $HW(s)$ models the different hardware components, such as the sensors, the actuators, the structural filters, etc. It is given by

$$HW(s) = \frac{0.000697s^2 - 0.0397s + 1}{0.000867s^2 + 0.0591s + 1}$$

We will denote by $G_0(s)$ the nominal plant and by $\hat{G}(s,X)$, (respectively $\hat{G}_0(s)$) the series connection $G(s,X)HW(s)$ (respectively $G_0(s)HW(s)$).

TABLE I

PARAMETERS FOR THE AIRCRAFT MODEL

Parameter	Mean	Standard Deviation
Z_α	-0.9381	0.0736
Z_q	0.0424	0.0035
M_α	1.6630	0.1385
M_q	-0.8120	0.0676
$Z_{\delta e}$	-0.3765	0.0314
$M_{\delta e}$	-10.8791	3.4695

We choose the controller to have the following structure

$$C(s,Y) = \begin{bmatrix} -K_a & -K_q\frac{(1+s\tau_1)}{(1+s\tau_2)} \end{bmatrix}$$

where the four parameters K_a, K_q, τ_1 and τ_2 have uniform distributions in the ranges

$$K_a \in [0,2],\ K_q \in [0,1],\ \tau_1 \in [0.01, 0.1],\ \tau_2 \in [0.01, 0.1].$$

We thus let $Y = [K_a \quad K_q \quad \tau_1 \quad \tau_2]^T$. Our objective is to find the controller which solves the following problem

$$\min \left\| W\left(I + \hat{G}C\right)^{-1} \right\|_\infty \text{ subject to } \left\| \frac{0.75C\hat{G}_0}{1+1.25C\hat{G}_0} \right\|_\infty \leq 1$$

where the weighting function $W(s)$ is given by

$$W(s) = \begin{bmatrix} \frac{2.8*6.28*31.4}{(s+6.28)(s+31.4)} & 0 \\ 0 & \frac{2.8*6.28*3.14}{(s+6.28)(s+31.4)} \end{bmatrix}$$

In order to adopt a randomized algorithm solution, in [68], this problem has been reformulated in the following way. Let us define a cost function

$$\Psi(Y) = \max\{\psi_1(Y), \psi_2(Y)\}$$

where

$$\psi_1(Y) = \begin{cases} 1 \text{ if } \left\| \frac{0.75C\hat{G}_0}{1+1.25C\hat{G}_0} \right\|_\infty > 1 \\ 0 \text{ otherwise} \end{cases}$$

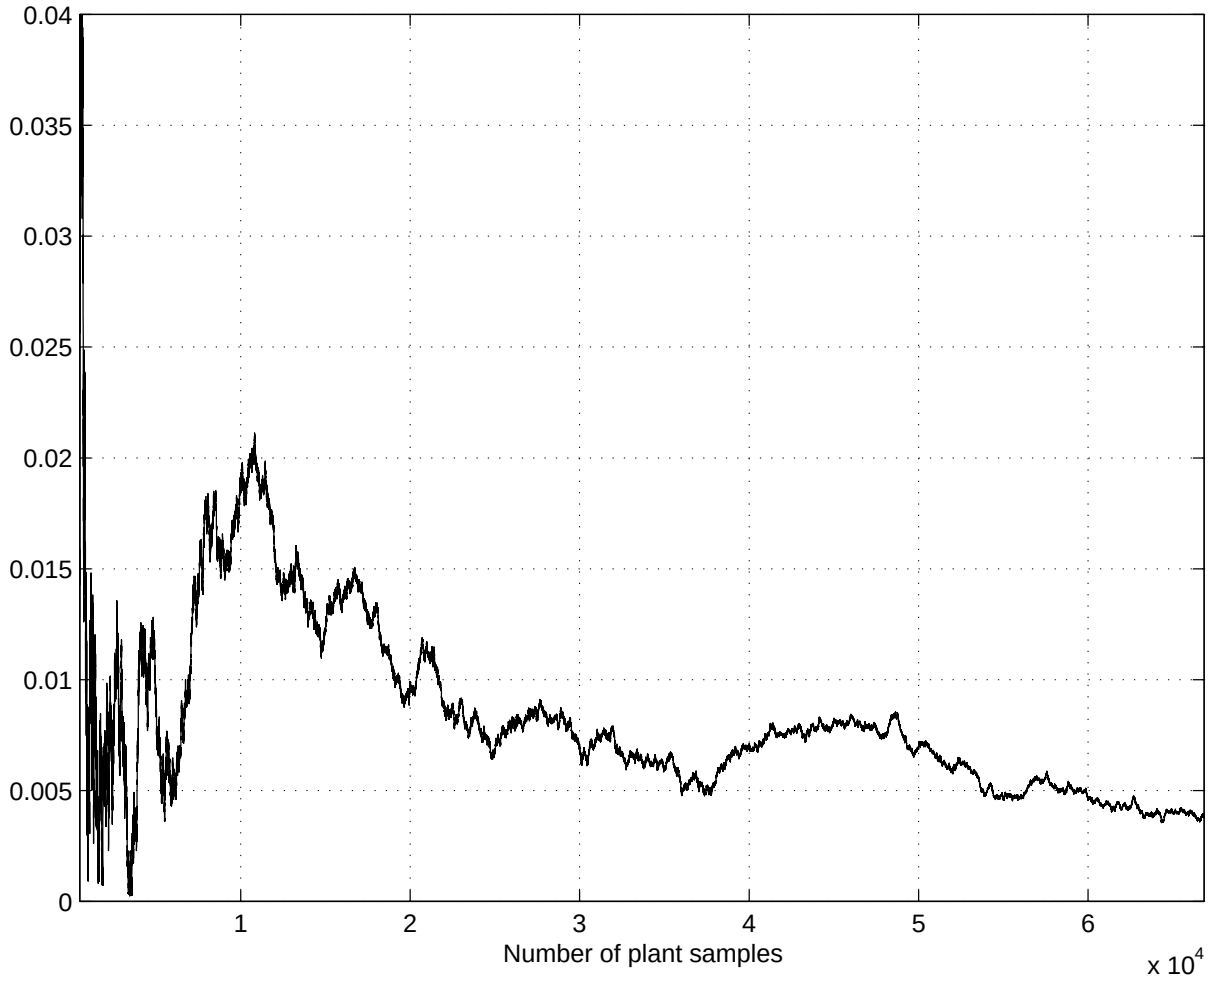

Fig. 3. The stopping variable

and

$$\psi_2(Y) = E_P\left(\zeta(X,Y)\right)$$

with

$$\zeta(X,Y) = \begin{cases} 1 \text{ if } (\hat{G}(X), C(Y)) \text{ is unstable} \\ \frac{\|W\left(I+\hat{G}(X)C(Y)\right)^{-1}\|_\infty}{1+\|W\left(I+\hat{G}(X)C(Y)\right)^{-1}\|_\infty} \text{ otherwise} \end{cases}$$

In our example, and for $\delta = 0.01$, $\alpha = 0.1$ and $\varepsilon = 0.1$, m evaluated to 51 controllers and n evaluated to $66,848$ plants and the procedure outlined in Algorithm 3 stopped after one iteration, i.e. $k = 1$. In Figure 3, the stopping variable is shown versus n. The parameters of the suboptimal

controller are

$$K_a = 1.7826,\ K_q = 0.7621,\ \tau_1 = 0.0511,\ \tau_2 = 0.0117,$$

and the corresponding value of the cost function is $\Psi(Y_{opt}) = 0.7149$, which compares favorably with the results of [68], where 2,619,047 plants were needed for the same ε, α, and δ.

△△△

Remark 2: As shown by Figure 3, the stopping condition is met far before the minimum numbers of plant samples n. This hints that with the same number of samples n, problems with much higher P-dimension could be addressed. The P-dimension of this problem was evaluated in [68] and is equal to $d = 118$. Therefore other types of controllers could be used instead of the first-order one we used. For instance the order of the controller could be increased until certain performance, in term of the desired value for Ψ_{opt}, is achieved [43].

◇◇◇

VI. Conclusions

In this paper we have drastically reduced the number of samples needed in order to obtain performance guarantees in robust control synthesis problems. This reduction is achieved by introducing sequential bootstrapping algorithms and exploiting the fact that the sample complexity is itself a random variable. This has allowed us to present Algorithm 3 as an efficient design methodology for fixed-order robust control design problems [43]. Recall for example that the Static Output Feedback (SOF) was shown in [14] to be NP-hard when the gains of the feedback matrix were bounded, but that Algorithm 3, is well suited to address exactly the SOF problem under those conditions.

It should be noted that the methodology presented in this paper can be used in many other application areas: one only needs to have an efficient analysis tool in order to convert it to an efficient design methodology. This is due to the fact that the design problem is converted to a sequence of analysis or verification problems after sampling more plants and controllers than the minimum number required by Algorithm 3. It should also be noted that the computational complexity or the undecidability of the problems studied are not eliminated but only avoided by relaxing the design requirements from absolute (hard) to probabilistic (soft) ones.

The randomized algorithms approach may be applied to design fixed-structure controllers for nonlinear systems (see for example the Pfaffian systems discussed in [66]), and to building software systems for practical control design problems. Our future research is concentrating at the theoretical level in obtaining better optimization algorithms and at the application level in designing software modules for linear and nonlinear control design. We are also investigating the applicability of the statistical learning approach in combination with the "unfalsified controller" design discussed in [54].

Acknowledgements: C.T. Abdallah acknowledges fruitful discussions with G.L. Heileman, and P. Goldberg.

Appendix

A. Proofs of Propositions and Theorems

Proof of Proposition 1. Let $S := \mathbb{R}^1$ and let $\mathcal{P}$ be the class of all symmetric continuous distributions on $\mathbb{R}^1$. Consider a class $\mathcal{F} := \{f\}$ consisting of only one function $f(x) = \text{sign}(x)$ (the range of the functions in this example is, of course, $[-1, 1]$, not $[0, 1]$). Then $Y_i = f(X_i)$ are i.i.d. with distribution $\mathbb{P}\{Y_i = 1\} = 1/2$ and $\mathbb{P}\{Y_i = -1\} = 1/2$ (for any $P \in \mathcal{P}$). Let the initial sample size $\bar{n}(\varepsilon, \delta)$ be equal to 1. Then

$$N(\varepsilon, \delta) = \min\{n \geq 1 : \mathbb{P}\{|\sum_{i=1}^{n} Y_i| \geq n\varepsilon\} \leq \delta\}.$$

$\{N(\varepsilon, \delta)\}$ is clearly a weakly efficient family of stopping times. We will prove that it is not a strongly efficient family. Let

$$\tau = \min\{n \geq 1 : \sum_{i=1}^{n} Y_i = 0\}.$$

First of all, the stopping time τ belongs to $\mathcal{T}_{\mathcal{F},\mathcal{P}}(\varepsilon, \delta)$ for all $\varepsilon, \delta > 0$. It is well known from Classical Probability that $\tau < +\infty$ with probability 1 (since the random walk $S_n := \sum_{j=1}^{n} Y_j$ is recurrent). Hence,

$$\mathbb{P}\{\tau > n\} \to 0 \text{ as } n \to \infty.$$

On the other hand, for $\varepsilon \in (0, 1)$,

$$\mathbb{P}\{|\sum_{i=1}^{n} Y_i| \geq n\varepsilon\} \geq \mathbb{P}\{Y_j = 1,\ j = 1, \ldots, n\} = 2^{-n}.$$

If we use the last bound with $n := N(\varepsilon; \delta)$, we get

$$\delta \geq \mathbb{P}\{|\sum_{i=1}^{n} Y_i| \geq n\varepsilon\} \geq \mathbb{P}\{Y_j = 1,\ j = 1, \ldots, n\} = 2^{-n},$$

which implies $N(\varepsilon; \delta) \geq \log_2(\delta^{-1})$. Hence,

$$\mathbb{P}\{\tau > N(\varepsilon; \delta)\} \to 0 \text{ as } \delta \to 0.$$

Suppose now that $\{N(\varepsilon, \delta)\}$ is strongly efficient. Then, for some $K_2 \geq 1, K_3 \geq 1$

$$1 - K_3\delta \leq \mathbb{P}\{\tau \geq N(K_2\varepsilon, \delta)\},$$

which contradicts the previous limit relationship. ∎

The proof of Proposition 2 follows from the concentration inequalities for empirical and Rademacher processes (see Lemmas 1, 2) and from a symmetrization inequality (see Lemma 3).

Lemma 1: For all $\varepsilon > 0$,

$$\mathbb{P}\{\|P_n - P\|_{\mathcal{F}} \geq \mathbb{E}\|P_n - P\|_{\mathcal{F}} + \varepsilon\} \leq \exp\{-\varepsilon^2 n/2\}$$

and

$$\mathbb{P}\{\mathbb{E}\|P_n - P\|_{\mathcal{F}} \geq \|P_n - P\|_{\mathcal{F}} + \varepsilon\} \leq \exp\{-\varepsilon^2 n/2\}.$$

■

Lemma 2: For all $\varepsilon > 0$,

$$\mathbb{P}\{\mathbb{E}\|n^{-1}\sum_{j=1}^{n} r_j \delta_{X_j}\|_{\mathcal{F}} \geq \|n^{-1}\sum_{j=1}^{n} r_j \delta_{X_j}\|_{\mathcal{F}} + \varepsilon\} \leq \exp\{-\varepsilon^2 n/4\}$$

and

$$\mathbb{P}\{\|n^{-1}\sum_{j=1}^{n} r_j \delta_{X_j}\|_{\mathcal{F}} \geq \mathbb{E}\|n^{-1}\sum_{j=1}^{n} r_j \delta_{X_j}\|_{\mathcal{F}} + \varepsilon\} \leq \exp\{-\varepsilon^2 n/4\}.$$

■

Lemma 3: The following inequality holds:

$$\frac{1}{2}\mathbb{E}\|n^{-1}\sum_{j=1}^{n} r_j(\delta_{X_j} - P)\|_{\mathcal{F}} \leq \mathbb{E}\|P_n - P\|_{\mathcal{F}} \leq 2\mathbb{E}\|n^{-1}\sum_{j=1}^{n} r_j \delta_{X_j}\|_{\mathcal{F}}.$$

■

The proofs of Lemmas 1,2 follow from the well known and widely used concentration inequalities for martingale difference sequences (see, e.g., Ledoux and Talagrand [44]), Lemma 1.5). See also [25], Theorems 9.1, 9.2. The proof of Lemma 3 can be found, for instance, in [62].

Proof of Proposition 2. Let $\bar{n} := \bar{n}(\varepsilon, \delta)$. Lemma 1 implies that

$$\mathbb{P}\Big(\bigcap_{n \geq \bar{n}} \Big\{\|P_n - P\|_{\mathcal{F}} \leq \mathbb{E}\|P_n - P\|_{\mathcal{F}} + \varepsilon\Big\}\Big) \geq 1 - \sum_{n \geq \bar{n}} \exp\{-\varepsilon^2 n/4\} \geq 1 - \frac{\exp\{-\varepsilon^2 \bar{n}/4\}}{1 - e^{-\varepsilon^2/4}} \geq 1 - \delta/2.$$

It follows from Lemma 3 that

$$\mathbb{P}\Big(\bigcap_{n \geq \bar{n}} \Big\{\|P_n - P\|_{\mathcal{F}} \leq 2\mathbb{E}\|n^{-1}\sum_{j=1}^{n} r_j \delta_{X_j}\|_{\mathcal{F}} + \varepsilon\Big\}\Big) \geq 1 - \delta/2.$$

Similarly, Lemma 2 implies that for all $\varepsilon > 0$

$$\mathbb{P}\Big(\bigcap_{n \geq \bar{n}} \Big\{ \mathbb{E}\|n^{-1} \sum_{j=1}^{n} r_j \delta_{X_j}\|_{\mathcal{F}} \leq \|n^{-1} \sum_{j=1}^{n} r_j \delta_{X_j}\|_{\mathcal{F}} + \varepsilon \Big\} \Big) \geq 1 - \delta/2.$$

Thus,

$$\mathbb{P}\Big(\bigcap_{n \geq \bar{n}} \Big\{ \|P_n - P\|_{\mathcal{F}} \leq 2\|n^{-1} \sum_{j=1}^{n} r_j \delta_{X_j}\|_{\mathcal{F}} + 3\varepsilon \Big\} \Big) \geq 1 - \delta.$$

Since for $n = \nu(\varepsilon, \delta)$ we have

$$\|n^{-1} \sum_{j=1}^{n} r_j \delta_{X_j}\|_{\mathcal{F}} \leq \varepsilon,$$

we get the first bound in proposition 2

$$\mathbb{P}\{\Delta_{\nu(\varepsilon,\delta)} \geq 5\varepsilon\} \leq \delta.$$

Quite similarly, it follows from Lemmas 1, 2 and 3 that for all $n \geq \bar{n}$ with probability $\geq 1 - \delta$

$$\begin{aligned}
\|P_n - P\|_{\mathcal{F}} &\geq \mathbb{E}\|P_n - P\|_{\mathcal{F}} - \varepsilon \\
&\geq \frac{1}{2}\mathbb{E}\|n^{-1} \sum_{j=1}^{n} r_j (\delta_{X_j} - P)\|_{\mathcal{F}} - \varepsilon \\
&\geq \frac{1}{2}\mathbb{E}\|n^{-1} \sum_{j=1}^{n} r_j \delta_{X_j}\|_{\mathcal{F}} - \frac{1}{2}\mathbb{E}|n^{-1} \sum_{j=1}^{n} r_j| - \varepsilon \\
&\geq \frac{1}{2}\mathbb{E}\|n^{-1} \sum_{j=1}^{n} r_j \delta_{X_j}\|_{\mathcal{F}} - \frac{1}{2}\mathbb{E}^{1/2}|n^{-1} \sum_{j=1}^{n} r_j|^2 - \varepsilon \\
&\geq \frac{1}{2}\|n^{-1} \sum_{j=1}^{n} r_j \delta_{X_j}\|_{\mathcal{F}} - \frac{1}{2\sqrt{n}} - \frac{3}{2}\varepsilon. \qquad (6)
\end{aligned}$$

Therefore, with probability $\geq 1 - \delta$, the condition $\bar{n} \leq n < \nu(6\varepsilon, \delta)$ (note also that $\bar{n} \geq \bar{n}(6\varepsilon; \delta)$) implies that

$$\|n^{-1} \sum_{j=1}^{n} r_j \delta_{X_j}\|_{\mathcal{F}} > 6\varepsilon$$

and, in view of the previous inequalities, we get (taking into account $n \geq \bar{n}(\varepsilon, \delta)$ and $\bar{n}(\varepsilon, \delta)^{-1/2} \leq \varepsilon$ for all $\delta \in (0, 1)$) that with probability $\geq 1 - \delta$

$$\Delta_n > \frac{3}{2}\varepsilon - \frac{1}{2\sqrt{\bar{n}(\varepsilon, \delta)}} \geq \varepsilon \text{ for all } \bar{n}(\varepsilon; \delta) \leq n < \nu(6\varepsilon; \delta).$$

The second bound in proposition 2 now easily follows. ■

The proof of Theorem 5 requires the following lemma (which can be proven along the same lines as, for instance, Lemma 2.3.7 in [62]).

Lemma 4: Suppose Z_1, Z_2 are independent stochastic processes in $\ell^\infty(\mathcal{F})$. Then for all $t > 0, c > 0$

$$\mathbb{P}\{\|Z_1\|_{\mathcal{F}} \geq t + c\} \leq \frac{\mathbb{P}\{\|Z_1 - Z_2\|_{\mathcal{F}} \geq t\}}{\inf_{f\in\mathcal{F}} \mathbb{P}\{|Z_2(f)| \leq c\}}.$$

■

Proof of Theorem 5. The proof of the first statement is quite similar to the first part of the proof of proposition 2. We set $\bar{n} := \bar{n}(\varepsilon;\delta)$. For instance, instead of the bound (4), we have here

$$\begin{aligned}\mathbb{P}\Big(\bigcap_{n\in\{2^k\bar{n}:k=0,1,\ldots\}} \Big\{\|P_n - P\|_{\mathcal{F}} \leq \mathbb{E}\|P_n - P\|_{\mathcal{F}} + \varepsilon\Big\}\Big) &\geq 1 - \sum_{k=0}^{\infty} \exp\{-\varepsilon^2\bar{n}2^k/4\} \\ &\geq 1 - 2\exp\{-\varepsilon^2\bar{n}/4\} \\ &\geq 1 - \delta/2\end{aligned}$$

where we have used the fact that for any $\alpha \geq 1$ we have

$$\begin{aligned}\sum_{k=1}^{\infty} \exp\{-\alpha(2^k - 1)\} &\leq \sum_{k=1}^{\infty} \exp\{-(2^k - 1)\} \\ &\leq \sum_{k=1}^{\infty} e^{-k} = (e-1)^{-1} \\ &< 1\end{aligned}$$

and hence

$$\sum_{k=0}^{\infty} \exp\{-\alpha 2^k\} \leq 2e^{-\alpha}.$$

Similar minor changes are needed in other parts of the proof.

To prove the second property in the definition of the weakly efficient stopping times, let $N := N(\varepsilon;\delta)$, let $n_k := 2^k\bar{n}(24\varepsilon;\delta)$ and choose k such that $n_k \leq N < n_{k+1}$. Then

$$\mathbb{P}\{\nu(24\varepsilon;\delta) > N\} \leq \mathbb{P}\{\nu(24\varepsilon;\delta) > n_k\}.$$

If $\nu(24\varepsilon;\delta) > n_k$, then for $n = n_k$

$$\|n^{-1}\sum_{j=1}^{n} r_j\delta_{X_j}\|_{\mathcal{F}} > 24\varepsilon.$$

Since, by the assumptions, $N \geq \bar{n}$, we get $n_k \geq \bar{n}/2$. Similarly to the proof of the bound (6), we obtain that with probability $\geq 1 - \delta$

$$\|P_n - P\|_{\mathcal{F}} \geq \frac{1}{2}\|n^{-1}\sum_{j=1}^{n} r_j\delta_{X_j}\|_{\mathcal{F}} - \frac{1}{2\sqrt{n}} - 6\varepsilon,$$

which implies that

$$\begin{aligned}\mathbb{P}\{\nu(24\varepsilon;\delta) > n_k\} &\leq \mathbb{P}\{\|P_{n_k} - P\|_{\mathcal{F}} \geq 4\varepsilon\} + \delta \\ &= \mathbb{P}\{\|S_{n_k}\|_{\mathcal{F}} \geq 4\varepsilon n_k\} + \delta \\ &\leq \mathbb{P}\{\|S_{n_k}\|_{\mathcal{F}} \geq 2\varepsilon N\} + \delta,\end{aligned}$$

where

$$S_n(f) := \sum_{j=1}^{n}[f(X_j) - P(f)],\ f \in \mathcal{F}.$$

Next we use Lemma 4,

$$\mathbb{P}\{\|S_{n_k}\|_{\mathcal{F}} \geq 2\varepsilon N\} \leq \frac{\mathbb{P}\{\|S_N\|_{\mathcal{F}} \geq \varepsilon N\}}{\inf_{f\in\mathcal{F}} \mathbb{P}\{|(S_N - S_{n_k})(f)| \leq \varepsilon N\}}.$$

and by Hoeffding's inequality [66]

$$\begin{aligned}\inf_{f\in\mathcal{F}} \mathbb{P}\{|(S_N - S_{n_k})(f)| \leq \varepsilon N\} &= 1 - \sup_{f\in\mathcal{F}} \mathbb{P}\{|(S_N - S_{n_k})(f)| > \varepsilon N\} \\ &\geq 1 - 2\exp\{-\varepsilon^2 N/2\} \\ &\geq 1 - \delta,\end{aligned}$$

and we get

$$\begin{aligned}\mathbb{P}\{\|S_{n_k}\|_{\mathcal{F}} \geq 2\varepsilon N\} &\leq (1-\delta)^{-1}\mathbb{P}\{\|P_N - P\|_{\mathcal{F}} \geq \varepsilon\} \\ &\leq \delta(1-\delta)^{-1}.\end{aligned}$$

Hence, we get

$$\begin{aligned}\mathbb{P}\{\nu(24\varepsilon;\delta) > n_k\} &\leq \delta(1-\delta)^{-1} + \delta \\ &\leq 3\delta,\end{aligned}$$

for $\delta < 1/4$, which implies weak efficiency with $K_1 = 5, K_2 = 24$ and $K_3 = 3$. ■

Proof of Proposition 3. By Hoffmann-Jorgensen inequality (see van der Vaart and Wellner [62]), it follows that

$$\mathbb{E}\|P_N - P\|_{\mathcal{F}} \leq 12F^{-1}(47/48) + 12\mathbb{E}\max_{1\leq n\leq N}\left\|\frac{f(X_k) - Pf}{N}\right\|_{\mathcal{F}},$$

where F^{-1} is the quantile function (the inverse of the distribution function) of $\|P_N - P\|_{\mathcal{F}}$. Since

$$\mathbb{P}\{\|P_N - P\|_{\mathcal{F}} \geq \varepsilon\} \leq \delta$$

and $\delta \leq 1/48$ we get $F^{-1}(47/48) \leq \varepsilon$. Hence,

$$\mathbb{E}\|P_N - P\|_{\mathcal{F}} \leq 12\varepsilon + 24/N \leq 14\varepsilon.$$

Define $n_k := 2^k \bar{n}(30\varepsilon; \delta)$. Choose k such that $n_{k-1} \leq \bar{n}(\varepsilon; \delta) < n_k$. Clearly, $n_k \leq 2\bar{n}(\varepsilon; \delta)$. By the submartingale property of the sequence $\|P_n - P\|_{\mathcal{F}}$ (see van der Vaart and Wellner [62]), for $n = n_k$, we have $\mathbb{E}\|P_n - P\|_{\mathcal{F}} \leq 14\varepsilon$, and by Lemma 3

$$\begin{aligned} \mathbb{E}\|n^{-1}\sum_{j=1}^{n} r_j \delta_{X_j}\|_{\mathcal{F}} &\leq 2\mathbb{E}\|P_n - P\|_{\mathcal{F}} + 2n^{-1/2} \\ &\leq 28\varepsilon + \varepsilon/4 \\ &\leq 29\varepsilon. \end{aligned}$$

By Lemma 2,

$$\begin{aligned} \mathbb{P}\{\|n^{-1}\sum_{j=1}^{n} r_j \delta_{X_j}\|_{\mathcal{F}} \geq 30\varepsilon\} &\leq \mathbb{P}\{\|n^{-1}\sum_{j=1}^{n} r_j \delta_{X_j}\|_{\mathcal{F}} \geq \mathbb{E}\|n^{-1}\sum_{j=1}^{n} r_j \delta_{X_j}\|_{\mathcal{F}} + \varepsilon\} \\ &\leq \frac{\delta}{4} \end{aligned}$$

which immediately implies the bound (5) with $K_1 = 30, K_2 = 2$. ■

References

[1] C. Abdallah, P. Dorato, W. Yang, R. Liska, and S. Steinberg. Application of quantifier elimination theory to control system design. In *Proceedings of the 4th IEEE Mediterranean Symposium on Control & Automation*, pages 41–44, Chania, Crete, Greece, 1996.

[2] H. Anai. On solving semidefinite programming by quantifier elimination. In *Proc. IEEE American Control Conf.*, pages 2814–2818, Philadelphia, PA, 1998.

[3] B. Anderson, N. Bose, and E. Jury. Output feedback stabilization and related problems – solution via decision methods. *IEEE Trans. Automat. Control*, 20:53–65, 1975.

[4] K. Aström and T. Hägglund. *Automatic Tuning of PID Controllers.* Instrument Society of America, Research Triangle Park, NC, 1988.

[5] K. Azuma. Weighted sums of certain dependent random variables. *Tokuku Math. J.*, 19:357–367, 1967.

[6] E. Bai, R. Tempo, and M. Fu. Worst case properties of the uniform distribution and randomized algorithms for robustness analysis. In *Proc. IEEE American Control Conf.*, pages 861–865, Albuquerque, NM, 1997.

[7] B. Barmish and C. Lagoa. The uniform distribution: A rigorous justification for its use in robustness analysis. In *Proc. IEEE Conf. on Dec. and Control*, pages 3418–3423, Kobe, Japan, 1996.

[8] B. Barmish and C. Lagoa. The uniform distribution: A rigorous justification for its use in robustness analysis. *Mathematics of Control, Signals, and Systems*, 10:203–222, 1997.

[9] B. Barmish, C. Lagoa, and R. Tempo. Radially truncated uniform distributions for probabilistic robustness of control systems. In *Proc. IEEE American Control Conf.*, pages 853–857, Albuquerque, NM, 1997.

[10] B. Barmish and R. Tempo. Probabilistic robustness: A new line of research. Tutorial Workshop, IEEE Conf. on Dec. and Control, San Diego, CA, 1997.

[11] S. Basu, R. Pollack, and M.-F. Roy. On the combinatorial and algebraic complexity of quantifier elimination. In *Proc. 35th Symposium on Foundations of Computer Science*, pages 632–641, Santa Fe, NM, November 22–24 1994.

[12] V. Blondel. *Simultaneous Stabilization of Linear Systems.* Springer-Verlag, London, 1st edition, 1994.

[13] V. Blondel, E. Sontag, M. Vidyasagar, and J. Willems, editors. *Simultaneous Stabilization of Linear Systems and Interpolation with Rational Functions.* Springer Verlag, London, 1999.

[14] V. Blondel and J. Tsitsiklis. NP-hardness of some linear control design problems. *SIAM J. of Control and Opt.*, 35(6):2118–2127, 1997.

[15] V. Blondel and J. Tsitsiklis. NP-hardness of some linear control design problems. *SIAM Journal of Control and Optimization*, 35:2118–2127, 1997.

[16] V. Blondel and J. Tsitsiklis. A survey of computational complexity results in systems and control. Preprint available at http://www.ulg.ac.be/mathsys/blondel/publications.html, 1998.

[17] V. Blondel and J. Tsitsiklis. Three problems on the decidability and complexity of stability. In V. Blondel, E. Sontag, M. Vidyasagar, and J. Willems, editors, *Open Problems in Mathematical Systems and Control Theory.* Springer Verlag, London, 1999.

[18] X. Chen and K. Zhou. On the probabilistic characterization of model uncertainty and robustness. In *Proc. IEEE Conf. on Dec. and Control*, pages 3816–3821, San Diego, CA, 1997.

[19] X. Chen and K. Zhou. A probabilistic approach for robust control. In *Proc. IEEE Conf. on Dec. and Control*, pages 4894–4895, San Diego, CA, 1997.

[20] X. Chen and K. Zhou. Constrained optimal synthesis and robustness analysis by randomized algorithms. In *Proc. IEEE American Control Conf.*, pages 1429–1433, Philadelphia, PA, 1998.

[21] H. Chernoff. A measure of asymptotic efficiency for tests of a hypothesis based on the sum of observations. *Ann. Math. Statist.*, pages 493–507, 1952.

[22] G. E. Collins. Quantifier elimination in the elementary theory of real closed fields by cylindrical algebraic decomposition. In *Lecture Notes in Computer Science*, Vol. 33, pages 134–183, Berlin, Springer-Verlag, Berlin, 1975.

[23] G. E. Collins and H. Hong. Partial cylindrical algebraic decomposition for quantifier elimination. *J. Symbolic Computation*, 12(3):299–328, 1991.

[24] G. Coxson. *Computational Complexity of Robust Stability and Regularity in Families of Linear Systems.* PhD thesis, The University of Wisconsin–Madison, 1993.

[25] L. Devroy, L. Györfi, and G. Lugosi. *A Probabilistic Theory of Pattern Recognition.* Springer-Verlag, Berlin, 1996.

[26] P. Dorato, W. Yang, and C. Abdallah. Robust Multi-Objective Feedback Design by Quantifier Elimination. *J. Symbolic Computation*, 24:153–159, 1997.

[27] J. C. Doyle and G. Stein. Multivariable feedback design: Concepts for a classical/modern synthesis. *IEEE Trans. Automat. Control*, 26:4–16, 1981.

[28] R. Dudley. Central limit theorems for empirical measures. *Annals of Probability*, 6:899–929, 1978.

[29] R. Dudley. *Uniform Central Limit Theorems.* In Press, Cambridge University Press, Cambridge, 1999.

[30] R. Durrett. *Probability: Theory and Examples.* 2nd Edition, Duxbury Press, Belmont, CA, 1996.

[31] R. Erwin, A. Sparks, and D. Bernstein. Fixed-structure robust controller synthesis via decentralized static output feebdack. *International Journal of Robust and Nonlinear Control*, 8:499–522, 1998.

[32] M. Garey and D. Johnson. *Computers and Intractability: A Guide to the Theory of NP-Completeness.* W.H. Freeman and Co., New York, N.Y., 1979.

[33] R. Hamming. *Coding and Information Theory.* Prentice-Hall, Inc., Englewood Cliffs, N.J., 1986.

[34] T. Hgerup and C. Rüb. A guided tour of Chernoff bounds. *Information Processing letters*, 33:305–308, 1989.

[35] H. Hong. *Improvements in CAD-based Quantifier Elimination.* PhD thesis, The Ohio State University, Columbus, OH, 1990.

[36] H. Hong. Simple solution formula construction in cylindrical algebraic decomposition based quantifier elimination. In P. S. Wang, editor, *ISSAC'92, International Symposium on Symbolic and Algebraic Computation, July 27-29, Berkeley, California*, pages 177–188, ACM Press, New York, 1992.

[37] M. Jirstrand. *Algebraic Methods for Modeling and Design in Control.* PhD thesis, Linköping University - Sweden, 1996. Thesis Nr. 540.

[38] M. Karpinski and A.J. Macintyre. Polynomial bounds for VC dimension of sigmoidal and general Pfaffian neural networks. *J. Comp. and Sys. Sci.*, 54(1):169-176, February 1997.

[39] L. Keel and S. Bhattacharyya. Robust, fragile, or optimal? *IEEE Trans. Autom. Control*, AC-42:1098–1105, 1997.

[40] P. Khargonekar and A. Tikku. Randomized algorithms for robust control analysis and synthesis have polynomial complexity. In *Proc. IEEE Conf. on Dec. and Control*, pages 3470–3475, Kobe, Japan, 1996.

[41] V. Koltchinskii. Functional limit theorems and empirical entropy, 1. *Teorija Verojatnosteh i Matematicheskaja Statistika*, 33:31–42, 1985. (English translation: *Probability Theory and Mathematical Statistics*, 33, 1986)

[42] V. Koltchinskii. Functional limit theorems and empirical entropy, 2. *Teorija Verojatnosteh i Matematicheskaja Statistika*, 34:73–85, 1986. (English translation: *Probability Theory and Mathematical Statistics*, 33, 1987)

[43] V. Koltchinskii, M. Ariola, C.T. Abdallah and P. Dorato. Statistical controller design for the linear benchmark problem. *Submitted, 1999 Conf. Dec. and Control*, Phoenix, AZ, 1999.

[44] M. Ledoux and M. Talagrand. *Probability in Banach spaces.* Springer-Verlag, Berlin, 1991.

[45] L.H. Lee and K. Poolla. Statistical validation for uncertainty models. In B. Francis and A.R. Tannenbaum, editors, *Feedback Control, Nonlinear Systems, and Complexity*, pages 131–149, Springer Verlag, London, 1995.

[46] C. McDiarmid. On the method of bounded differences. In *Surveys in Combinatorics*, pages 148–188. Cambridge University Press, Cambridge, 1989.

[47] V. Milman and G. Schechtman. *Asymptotic theory of finite dimensional normed spaces.* LNM, 1200, Springer-Verlag, New York, 1986.

[48] M. Morari and E. Zafiriou. *Robust Process Control.* Prentice-Hall, Upper Saddle River, NJ, 1989.

[49] A. Nemirovskii. Several NP-hard problems arising in robust stability analysis. *Mathematics of Control, Signals, and Systems*, 6:99–105, 1993.

[50] D. Neši'c and I. Mareels. Dead beat controllability of polynomial systems: Symbolic computation approaches. *IEEE Trans. Automat. Control*, 43:162–175, 1998.

[51] S. Poljak and J. Rohn. Checking robust nonsigularity is NP-hard. *Mathematics of Control, Signals, and Systems*, 6:1–9, 1993.

[52] L. Ray and R. Stengel. A Monte Carlo approach to the analysis of control system robustness. *Automatica*, 29:229–236, 1993.

[53] W. Rhee and M. Talagrand. Martingale inequalities and NP-complete problems. *Math. Operation Research*, 12:177–181, 1987.

[54] M. Safonov and T-C Tsao. The unfalsified control concept: a direct path from experiment to controller. In B. Francis and A.R. Tannenbaum, editors, *Feedback Control, Nonlinear Systems, and Complexity*, pages 196–214, Springer Verlag, London, 1995.

[55] R. Stengel and L. Ray. Stochastic robustness of linear time-invariant control systems. *IEEE Trans. Automat. Control*, 36:82–87, 1991.

[56] V. Syrmos, C. Abdallah, P. Dorato, and K. Grigoriadis. Static output feedback: A survey. *Automatica*, 33(2):125–1379, 1997.

[57] M. Talagrand. A new look at independence. *Annals of Probability*, 24:1–34, 1996.

[58] A. Tarski. *A Decision Method for Elementary Algebra and Geometry.* Univ. of California Press, Berkeley, 2nd Ed, 1951.

[59] R. Tempo, E. Bai, and F. Dabbene. Probabilistic robustness analysis: Explicit bounds for the minimum number of samples. *Systems and Control Letters*, 30:237–242, 1997.

[60] R. Tempo and F. Dabbene. Probabilistic robustness analysis and design of uncertain systems. Preprint, 1998.

[61] O. Toker and H. Özbay. On the NP-hardness of solving bilinear matrix inequalities and simultaneous stabilization with static output feedback. *Proc. IEEE American Control Conf.*, pages 2056–2064, Seattle, WA, June 1995.

[62] A. van der Vaart and J. Wellner. *Weak Convergence and Empirical Processes*. Springer-Verlag, Berlin, 1996.

[63] V. Vapnik. *Estimation of Dependences Based on Empirical Data*. Springer-Verlag, New York, NY, 1982.

[64] V. Vapnik. *Statistical Learning Theory*. John Wiley & Sons, New York, NY, 1998.

[65] V. Vapnik and A. Chervonenkis. On the uniform convergence of relative frequencies of events to their probabilities. *Theory of Probability and its Applications*, 16:264–280, 1971.

[66] M. Vidyasagar. *A Theory of Learning and Generalization with Applications to Neural Networks and Control Systems*. Springer-Verlag, Berlin, 1996.

[67] M. Vidyasagar. Statistical learning theory and its applications to randomized algorithms for robust controller synthesis. In *European Control Conference (ECC97)*, Plenary Lectures and Mini-Courses Volume, pages 162–190, Brussels, Belgium, 1997.

[68] M. Vidyasagar. Statistical learning theory and randomized algorithms for control. *IEEE Control Systems Magazine*, 18(6):69–85, 1998.

[69] V.V. Yurinski. Exponential bounds for large deviations. *Theory Prob. appl.*, 19:154–155, 1974.

[70] X. Zhu, Y. Huang, and J. Doyle. Soft vs hard bounds in probabilistic robustness analysis. In *Proc. IEEE Conf. on Dec. and Control*, pages 3412–3417, Kobe, Japan, 1996.