

OPTIMAL ASYMPTOTIC IDENTIFICATION UNDER BOUNDED DISTURBANCES

David N.C. Tse¹
Munther A. Dahleh²
John N. Tsitsiklis³

Laboratory for Information and Decision Systems, M.I.T.
Cambridge, MA 02139

ABSTRACT

This paper investigates the intrinsic limitation of identification of LTI systems using data corrupted by bounded disturbances, when the unknown plant is known to belong to a given *model set*. This is done by analyzing the optimal worst-case asymptotic error achievable by performing experiments using *any* bounded inputs and estimating the plant using *any* identification algorithm. The solution involves two steps. First, it is shown that under some compactness conditions on the model set, the asymptotic problem is equivalent to an *infinite-horizon* problem, where the entire infinite data record is available to compute the estimate. Second, the infinite-horizon problem is analyzed to find optimal inputs for minimizing the worst-case error. Using the l_1 norm as the error criterion, an interesting dichotomy of the class of all balanced and convex model sets is obtained: either inputs can be chosen to identify the plant with very small asymptotic error, or no finite set of inputs can yield finite error. Applying the general results on specific model sets, it is shown that the class of all stable plants and the class of all finite-dimensional plants belong to the first category of the dichotomy. Explicit characterization of the optimal inputs for these model sets are also obtained.

¹Research supported by an NSERC fellowship from the government of Canada, and by the NSF under Grant ECS-8552419

²Research supported by the ARO under Grant DAAL03-86-K-0171

³Research supported by the NSF under Grant ECS-8552419 and by the ARO under Grant DAAL03-86-K-0171

1 Introduction

This paper addresses some fundamental issues that arise in the field of system identification. These issues are expected to have important ramifications for robust and adaptive control.

Our work is motivated by the following questions: Given some *a priori* knowledge about an unknown plant and the measurement noise, together with some constraints on the experiments that can be performed, is it possible to identify the plant accurately enough so that the resulting model can be used for a prescribed objective? More generally, how much *a priori* knowledge is necessary for accurate identification, where accuracy is judged in relation to a specific purpose? And how is the achievable accuracy improved with additional *a priori* information?

Although these questions are fundamental, they have received little attention in the literature. While much of the previous work has been concerned with construction of efficient identification algorithms for different problems, our intention here is to study whether there *exists* an algorithm at all which yields the desired amount of identification, given the *a priori* knowledge available. The answer to this question depends only on the nature of the *a priori* knowledge and the experiments that one is allowed to perform, but is completely independent of any identification algorithm.

The questions we have raised are very general and, for the purpose of this paper, it is necessary to restrict the scope. We will deal exclusively with discrete-time, single-input-single-output linear time-invariant systems. The *a priori* information about the unknown plant is the knowledge that it belongs to a certain subset of the space of all LTI systems; this subset will be called a *model set*. The plant will be identified using input-output experiments. The inputs can be chosen freely but subject to a boundedness constraint. The observed output is corrupted by an additive disturbance (noise) which is assumed to be bounded but otherwise unknown. This is less restrictive than the usual stochastic noise assumptions, but while a probabilistic model allows an average case analysis of the identification error, the boundedness assumption forces us to consider worst-case errors since the disturbances can be arbitrarily adverse to the identification procedure. This situation can be viewed as a game played between the experimenter and an omnipotent adversary who attempts to choose a disturbance to ruin the accuracy of the estimate. We seek results on the optimal worst-case *asymptotic* error achievable by the experimenter when he is allowed to select the inputs, perform experiments of arbitrarily long length, and use the output measurements to generate an estimate, via whatever identification algorithm. Our goal is to investigate the key properties of model sets which can be identified with a small optimal error, and in particular how large the model set can be to still yield a finite optimal error. Furthermore, we are interested in robustness issues: does the optimal error vanish as the bound on the output disturbance decreases to zero? Answers to these

questions give a characterization of the difficulty of identification using a given model set.

The criterion used to measure the identification error depends on the intended use of the identified model. In this paper, we shall use the l_1 norm to measure the distance of the true impulse response from the identified one. Because the l_1 norm is an induced operator norm, it measures the quality of input/output predictions that can be made on the basis of the identified model. Moreover, this error measure is particularly suitable for applying certain robust control methodologies on the identified plant [2, 3, 8].

We will now briefly review the relevant literature. Some of the questions raised here are dealt with in a more general context by a *theory of optimal algorithms* [14, 22, 23], also known as *information-based complexity*. This theory provides a general mathematical framework for analyzing the optimal error achievable in solving a problem using a given amount of possibly inaccurate and partial information. *Information* plays the central role in this theory: the results depend only on the information used by an algorithm but are independent of its structure. Some of the basic ideas of this theory have influenced our work. Our research is also motivated by the “theory of the learnable” developed in the computer science field, following the pioneering work of Valiant [24]. Although the problems we consider are different, there is a similar emphasis on the intrinsic (algorithm-independent) difficulty of learning or identification problems. The problem of choosing good inputs for identification, or *experiment design*, has been subjected to wide attention in the literature. However, most of the work is done in a stochastic and parametric setting [5, 13, 27], and input optimization is performed by minimizing a scalar functional of the Cramer-Rao bound on the covariance of the parameter estimates. Outside mainstream stochastic identification research, there is a line of work which deals with identification under bounded disturbances [4, 10, 15, 16, 17, 19, 21]. All of this work deals with models of known order and the emphasis is on deriving efficient recursive algorithms for estimating the parameters, not on analyzing the optimal worst-case error achievable in the limit. Although Fogel and Huang [4] obtained some conditions for the convergence of the estimates, these conditions imposed certain requirements on the disturbances and hence their work is not in the worst-case spirit. Very recently, there has been some work which proposes that system identification should be done with the requirements of robust controller design in mind. From this viewpoint, a system identification method should not only produce an identified model but also a measure of its inaccuracy in a metric which is consistent with robust controller design. Helmicki et al [6, 7] considered this problem and the metric they employed is the H_∞ norm. However, they only dealt with one specific model set: a set of stable LTI systems for which a lower bound on the decay rate and an upper bound on the gain are known. They did not study whether these conditions are necessary for accurate identification.

The contributions of this paper are two-fold.

At a more general level, it introduces a framework for the analysis of optimal worst-case asymptotic error under bounded disturbances. The central result here is that, under some topological conditions on the model set, *infinite-horizon* experiments, where the entire infinite data record is available to compute estimates, can be viewed as a limit of *finite-horizon* experiments, where only finite data records are available. Analysis of optimal asymptotic error is then reduced to finding optimal inputs to minimize the worst-case error for the infinite-horizon problem. Although the framework is presented here in terms of the l_1 error measure, it is equally suitable for analyzing problems with different error norms, and many of the results reported here can be generalized for arbitrary norms.

At a more specific level, concrete results are obtained for optimal l_1 identification, when both the inputs and the disturbances are constrained to be magnitude-bounded. It is shown that a lower bound to the optimal asymptotic error for *any* non-trivial model set is given by δ , the bound on the disturbance. For the important class of balanced convex model sets, an interesting dichotomy is proven: that either inputs can be chosen to identify the plant very accurately with an asymptotic error of less than 2δ , or no finite set of inputs can yield finite asymptotic error. Moreover, it is shown that large model sets, such as the class of all stable plants and the class of all finite-dimensional plants, belong to the first category of the dichotomy. Explicit characterization of the optimal inputs is also obtained.

The organization of the paper is as follows. In Section 2, the identification problem is formulated and the optimal worst-case asymptotic error achievable by any identification algorithm is defined. In Section 3, we present consistency results establishing infinite-horizon experiments as limits of finite-horizon ones. General results, giving conditions for inputs to be optimal for the infinite-horizon problem, are given in Section 4. In Section 5, the general results developed are applied to analyse specific model sets. In Section 6, we explore the question of when the convergence of the finite-horizon identification errors to the infinite-horizon error is uniform over all unknown plants in the model set. In Section 7, we show that for balanced convex model, the use of feedback (using observed outputs to decide future inputs) cannot lead to significant improvement in the identification error. Section 8 contains our conclusions, and discusses future research direction and in particular the implications of this work to robust and adaptive control.

2 Problem Formulation

Let \mathcal{X} be the class of all causal, single-input single-output, linear time-invariant, discrete-time systems. We identify \mathcal{X} with the space of all one-sided real-valued sequences, \mathfrak{R}^ω . Let $\mathcal{M} \subset \mathcal{X}$ be the *model set* which is assumed to contain the unknown plant h to be identified. The set \mathcal{M} captures the experimenter's *a priori* knowledge about h . Also given is an input set \mathcal{U} which contains all the input sequences that can be used in the identification experiments. In this paper, we will take \mathcal{U} to be $\overline{BI}_\infty \equiv \{u : \|u\|_\infty \leq 1\}$, where $\|u\|_\infty \equiv \sup_i |u_i|$. Indeed, it is reasonable to assume that unbounded signals should not be inputs to the system, either during the identification or the operation stage, because of physical limitations, power restrictions, safety, or in order to maintain the validity of the linear model of the plant.

An *experiment* is conducted by choosing an input sequence $u \in \overline{BI}_\infty$ and measuring the output sequence y , related to u by

$$y = h * u + d \quad (2.1)$$

where $*$ denotes the convolution operator and d is the disturbance sequence which corrupts the measurements. (Note that h, u, y, d are all one-sided real-valued sequences; $h = (h_0, h_1, h_2, \dots)$ etc.) The disturbance d is assumed to be magnitude bounded with $\|d\|_\infty \leq \delta$ for some known δ , but can otherwise be arbitrary. The disturbance may arise from actual measurement noise, such as quantization, or it may reflect non-linearities and time-variation of the plant. In the latter case, the true plant is actually nonlinear and time varying (as are most plants in reality) but is assumed to be approximated well at the operating range by an LTI component, which is the object of identification.

If N such experiments are performed, we have:

$$y^{(i)} = u^{(i)} * h + d^{(i)}, \quad i = 1, 2, \dots, N \quad (2.2)$$

where $y^{(i)}$ and $d^{(i)}$ are the output and disturbance sequences in the i th experiment. This can be written in a more compact notation:

$$\mathbf{y} = \mathbf{u} * h + \mathbf{d} \quad \|\mathbf{d}\|_\infty \equiv \max_i \|d^{(i)}\|_\infty \leq \delta \quad (2.3)$$

where $\mathbf{y} = [y^{(1)}, \dots, y^{(N)}]$, $\mathbf{u} = [u^{(1)}, \dots, u^{(N)}]$, and $\mathbf{d} = [d^{(1)}, \dots, d^{(N)}]$ are vectors of sequences; convolution of h with a vector of inputs is just element-wise convolution with every input. Also note that the vector of inputs \mathbf{u} is in \overline{BI}_∞^N .

An *identification algorithm* is a mapping ϕ which generates, at each time instant n , an estimate $\hat{h}^{(n)} \equiv \phi(P_n \mathbf{u}, P_n \mathbf{y}) \in \mathcal{X}$ of the unknown plant h , given the input and output sequences in the

experiments . Here, P_n is the truncation operator, defined by $P_n x = (x_0, x_1, \dots, x_n)$ for each infinite sequence x . Its use signifies that the algorithm ϕ generates at each time instant an estimate based only on the input-output data it has seen so far. In computing the estimates, the algorithm has access to what the model set \mathcal{M} is and also the value of δ , the bound on the disturbance. However, the estimates $\hat{h}^{(n)}$ are not required to belong to \mathcal{M} . This added generality may be of benefit in some circumstances.

To evaluate the accuracy of $\hat{h}^{(n)}$ as an estimate of h , the l_1 norm will be used, as defined by $\|g\|_1 \equiv \sum_{k=1}^{\infty} |g_k|$ for any plant g . The important property of the l_1 norm which makes it an appropriate error measure to use here is that it is an induced operator norm on the space of all magnitude-bounded signals:

$$\|g\|_1 = \sup_{\|u\|_{\infty} \leq 1} \|u * g\|_{\infty} \quad (2.4)$$

Thus, the error of an estimate measured in the l_1 norm can be interpreted as the *worst-case* prediction error when using the estimate to predict the output of the plant to *any* input in the input set $\mathcal{U} \equiv \overline{B}_1$.

Given an identification algorithm and a chosen set of input sequences for the experiments, we would like to consider the limiting situation when more and more data are observed. To this end, the worst case asymptotic error is defined as follows.

Definition 2.1 Fix the inputs \mathbf{u} . The worst-case asymptotic l_1 error, $e_{\infty}^1(\phi, \mathcal{M}, \mathbf{u}, \delta)$, of an algorithm ϕ is the smallest number r such that for all plants $h \in \mathcal{M}$ and for all disturbances \mathbf{d} with $\|\mathbf{d}\|_{\infty} \leq \delta$,

$$\limsup_{n \rightarrow \infty} \|\phi(P_n \mathbf{u}, P_n(\mathbf{u} * h + \mathbf{d})) - h\|_1 \leq r$$

Equivalently,

$$e_{\infty}^1(\phi, \mathcal{M}, \mathbf{u}, \delta) \equiv \sup_{h \in \mathcal{M}} \sup_{\|\mathbf{d}\|_{\infty} \leq \delta} \limsup_{n \rightarrow \infty} \|\phi(P_n \mathbf{u}, P_n(\mathbf{u} * h + \mathbf{d})) - h\|_1$$

According to this definition, no matter what the true plant and the disturbances are, the plant can be eventually approximated to within $e_{\infty}^1(\phi, \mathcal{M}, \mathbf{u}, \delta)$, using the estimates generated by the identification algorithm. This is quite analogous to the notion of convergence of estimates to the true plant in the classical probabilistic framework of identification. However, since the disturbances here are assumed to be arbitrary and not necessarily stationary, such convergence is not possible in general. Instead, we only require the estimates to enter and stay within a ball around the true plant rather than to converge to the exact plant.

The optimal worst-case asymptotic l_1 error $E_\infty^1(\mathbf{u}, \mathcal{M}, \delta)$ is defined as the smallest error achievable by any algorithm:

$$E_\infty^1(\mathbf{u}, \mathcal{M}, \delta) \equiv \inf_{\phi} e_\infty^1(\phi, \mathcal{M}, \mathbf{u}, \delta)$$

Any algorithm for which the infimum is attained is said to be *asymptotically optimal*. Furthermore, the optimal asymptotic error achievable by any input and any identification algorithm is defined as

$$E_\infty^1(\mathcal{M}, \delta) \equiv \inf_{\mathbf{u} \in \overline{Bl}_\infty^N} E_\infty^1(\mathbf{u}, \mathcal{M}, \delta)$$

The quantity $E_\infty^1(\mathcal{M}, \delta)$ can be viewed as the value of a game played between the experimenter and an adversary, with the experimenter selecting the best inputs and the best algorithm to minimize the asymptotic error of the estimates, and the adversary choosing the worst plant and the worst disturbance to maximize it. Thus, $E_\infty^1(\mathcal{M}, \delta)$ can be regarded as an intrinsic limitation to the identification accuracy, given the *a priori* knowledge about the plant (model set \mathcal{M}).

While the above definition of asymptotic error in the l_1 norm makes sense for stable plants, it is unsuitable for unstable ones. For example, the locations of the unstable poles of a plant will have to be identified *exactly* in order to have finite error in the l_1 norm, and this is an unrealistic requirement. The problem is that the l_1 norm measures error on the *entire* impulse response, which grows exponentially in the case of unstable plants. A more reasonable error measure in this case is to consider the l_1 norm only on *truncated* portions of the impulse response, i.e. $\|P_m g\|_1$. That is, we will compare the truncated estimate $P_m \phi(P_n \mathbf{u}, P_n(\mathbf{u} * h + \mathbf{d}))$ with the truncated impulse response. This leads to the definition of the following weaker notion of asymptotic error.

Definition 2.2 *The worst-case asymptotic truncated l_1 error, $e_\infty^t(\phi, \mathcal{M}, \mathbf{u}, \delta)$, of an algorithm ϕ on inputs \mathbf{u} is the smallest number r such that for all plants $h \in \mathcal{M}$ and for all disturbances \mathbf{d} with $\|\mathbf{d}\|_\infty \leq \delta$,*

$$\limsup_{n \rightarrow \infty} \|P_m \{\phi(P_n \mathbf{u}, P_n(\mathbf{u} * h + \mathbf{d})) - h\}\|_1 \leq r \quad \forall m$$

or equivalently,

$$e_\infty^t(\phi, \mathcal{M}, \mathbf{u}, \delta) \equiv \sup_{h \in \mathcal{M}} \sup_{\|\mathbf{d}\|_\infty \leq \delta} \sup_m \limsup_{n \rightarrow \infty} \|P_m \{\phi(P_n \mathbf{u}, P_n(\mathbf{u} * h + \mathbf{d})) - h\}\|_1$$

Whereas in the l_1 norm case we require that the estimates will eventually approximate the *entire* impulse response of the true plant, here we only require that the estimates eventually approximate *any truncated* portion of the true impulse response, although it may take longer to generate good estimates of a larger portion. The truncated l_1 norm, like the l_1 norm, measures the worst-case

error of using the estimate to predict the output of the true plant, but only for a finite duration. We can also define similar notions of optimal asymptotic truncated l_1 errors:

$$E_\infty^t(\mathbf{u}, \mathcal{M}, \delta) \equiv \inf_{\phi} e_\infty^t(\phi, \mathcal{M}, \mathbf{u}, \delta) \quad \text{and} \quad E_\infty^t(\mathcal{M}, \delta) \equiv \inf_{\mathbf{u} \in \overline{Bl}_\infty^N} E_\infty^t(\mathbf{u}, \mathcal{M}, \delta)$$

The following relationship between the two asymptotic errors is clear.

Proposition 2.3 *For all model set \mathcal{M} , inputs \mathbf{u} , algorithms ϕ and $\delta \geq 0$,*

$$e_\infty^t(\phi, \mathcal{M}, \mathbf{u}, \delta) \leq e_\infty^1(\phi, \mathcal{M}, \mathbf{u}, \delta)$$

We will now give a basic result about the worst-case asymptotic errors which will be useful for our later analysis.

Proposition 2.4 *For any model set \mathcal{M} , inputs $\mathbf{u} \in \overline{Bl}_\infty^N$, algorithms ϕ and $\delta \geq 0$,*

$$e_\infty^1(\phi, \mathbf{u}, \overline{\mathcal{M}}, \delta) \leq \lim_{x \downarrow \delta} e_\infty^1(\phi, \mathbf{u}, \mathcal{M}, x)$$

and

$$e_\infty^t(\phi, \mathbf{u}, \overline{\mathcal{M}}, \delta) \leq \lim_{x \downarrow \delta} e_\infty^t(\phi, \mathbf{u}, \mathcal{M}, x)$$

where $\overline{\mathcal{M}}$ is the closure of \mathcal{M} with respect to the l_1 -topology on \mathcal{X}

Proof. We shall only prove the result for the case of the asymptotic l_1 error. The proof for the truncated l_1 error follows similarly.

By definition, for all $x \geq 0$, and $\forall h \in \mathcal{M}$ and \mathbf{d} with $\|\mathbf{d}\|_\infty \leq x$, we have

$$\limsup_{n \rightarrow \infty} \|\phi(P_n \mathbf{u}, P_n(\mathbf{u} * h + \mathbf{d})) - h\|_1 \leq e_\infty^1(\phi, \mathbf{u}, \mathcal{M}, x) \quad (2.5)$$

Let $\epsilon > 0$. Take any $h \in \overline{\mathcal{M}}$ and $\|\mathbf{d}\|_\infty \leq \delta$. There exists a $h' \in \mathcal{M}$ such that $\|h - h'\|_1 \leq \epsilon$. Therefore

$$\limsup_{n \rightarrow \infty} \|\phi(P_n \mathbf{u}, P_n(\mathbf{u} * h + \mathbf{d})) - h\|_1 \quad (2.6)$$

$$\leq \limsup_{n \rightarrow \infty} \|\phi(P_n \mathbf{u}, P_n(\mathbf{u} * h' + \mathbf{u} * (h - h') + \mathbf{d})) - h'\|_1 + \epsilon \quad (2.7)$$

Now, $\|\mathbf{u} * (h - h') + \mathbf{d}\|_\infty \leq \delta + \epsilon$, so applying inequality (2.5) with $x = \delta + \epsilon$,

$$\limsup_{n \rightarrow \infty} \|\phi(P_n \mathbf{u}, P_n(\mathbf{u} * h' + \mathbf{u} * (h - h') + \mathbf{d})) - h'\|_1 \leq e_\infty^1(\phi, \mathcal{M}, \mathbf{u}, \delta + \epsilon) \quad (2.8)$$

It follows that

$$\limsup_{n \rightarrow \infty} \|\phi(P_n \mathbf{u}, P_n(\mathbf{u} * h + \mathbf{d})) - h\|_1 \leq e_\infty^1(\phi, \mathcal{M}, \mathbf{u}, \delta + \epsilon) + \epsilon \quad (2.9)$$

Letting ϵ go to 0 gives the desired result. \square

The objective of this paper is to analyze the optimal asymptotic errors for general model sets and then apply the results to specific ones. We will proceed in two stages:

1) Characterize the asymptotically optimal algorithms and characterize the resulting optimal worst-case asymptotic errors, in terms of the chosen inputs \mathbf{u} .

2) Given the characterization of the optimal asymptotic errors, determine the optimal selection of inputs \mathbf{u} which will make the asymptotic errors small.

3 Asymptotically Optimal Identification

The characterization of asymptotically optimal algorithms and optimal worst-case asymptotic errors is in terms of the important notion of the *uncertainty set*.

Definition 3.1 Let \mathbf{u} and \mathbf{y} be the input and measured output sequences, and δ be the bound on the disturbances. The finite-horizon uncertainty set at time n is defined to be

$$S_n(\mathcal{M}, \mathbf{u}, \mathbf{y}, \delta) = \{g \in \mathcal{M} : \|P_n(\mathbf{u} * g - \mathbf{y})\|_\infty \leq \delta\}$$

and the infinite-horizon uncertainty set is

$$S_\infty(\mathcal{M}, \mathbf{u}, \mathbf{y}, \delta) = \{g \in \mathcal{M} : \|\mathbf{u} * g - \mathbf{y}\|_\infty \leq \delta\}$$

The set S_n contains all the plants in the model set consistent with the output data seen until time n . It characterizes the uncertainty at time n : any plant in S_n can be the actual plant from the experimenter's point of view. Similarly, S_∞ contains all the plants that are consistent with the entire output sequences. It measures the uncertainty that the experimenter would still have even if he could perform infinitely long experiments and could see the entire output record. It is easy to see that the finite-horizon uncertainty sets become smaller with increasing n and shrink to S_∞ in the limit:

$$S_\infty(\mathcal{M}, \mathbf{u}, \mathbf{y}, \delta) = \bigcap_{n=0}^{\infty} S_n(\mathcal{M}, \mathbf{u}, \mathbf{y}, \delta)$$

For any set $A \subset \mathcal{X}$, define the diameter of the set A as

$$\text{diam}(A) = \sup_{g, h \in A} \|g - h\|_1$$

(The diameter of a set can be infinite.) We shall now define an important quantity.

Definition 3.2 Given a choice of the inputs \mathbf{u} , define the worst-case infinite-horizon diameter $D(\mathbf{u}, \mathcal{M}, \delta)$ to be the diameter of the largest possible uncertainty set :

$$D(\mathbf{u}, \mathcal{M}, \delta) \equiv \sup_{h \in \mathcal{M}} \sup_{\|\mathbf{d}\|_\infty \leq \delta} \text{diam}(S_\infty(\mathcal{M}, \mathbf{u}, \mathbf{u} * h + \mathbf{d}, \delta))$$

The quantity $D(\mathbf{u}, \mathcal{M}, \delta)$ is a measure of the worst uncertainty over all the possible output sequences that can be generated by any plant in the model set and any permissible disturbance. It turns out that it is precisely this quantity that characterizes the optimal worst-case asymptotic errors. First we show that half the worst-case infinite-horizon diameter is a lower bound to both optimal asymptotic errors.

Proposition 3.3 Let \mathcal{M} be any model set, \mathbf{u} be any vector of inputs and $\delta \geq 0$. Then

- (1) $e_\infty^t(\phi, \mathcal{M}, \mathbf{u}, \delta) \geq D(\mathbf{u}, \mathcal{M}, \delta)/2$ for any algorithm ϕ ; and
- (2) $e_\infty^1(\phi, \mathcal{M}, \mathbf{u}, \delta) \geq D(\mathbf{u}, \mathcal{M}, \delta)/2$ for any algorithm ϕ ;

Proof. Part (2) follows from part (1) in view of Proposition 2.3, so we only need to prove part (1). We shall assume $D(\mathbf{u}, \mathcal{M}, \delta) < \infty$. The case where $D(\mathbf{u}, \mathcal{M}, \delta)$ is infinite follows similarly.

For any $\epsilon > 0$, there exist two plants g and g' and an infinite-horizon uncertainty set $S_\infty(\mathcal{M}, \mathbf{u}, \mathbf{y}, \delta)$ such that $g, g' \in S_\infty(\mathcal{M}, \mathbf{u}, \mathbf{y}, \delta)$ and $\|g - g'\|_1 > D(\mathbf{u}, \mathcal{M}, \delta) - \epsilon$. We can pick integer $M > 0$ such that

$$\|P_M(g - g')\|_1 > D(\mathbf{u}, \mathcal{M}, \delta) - 2\epsilon \quad (3.10)$$

Also since both g and g' are in the same uncertainty set, there exist disturbances \mathbf{d} and \mathbf{d}' such that $\mathbf{u} * g + \mathbf{d} = \mathbf{u} * g' + \mathbf{d}' = \mathbf{y}$.

Let ϕ be any identification algorithm. Based on the output observation, ϕ cannot distinguish between g and g' . Hence, for every $m > M$ and for every n ,

$$\begin{aligned} & \|P_m(\phi(P_n \mathbf{u}, P_n(\mathbf{u} * g + \mathbf{d})) - g)\|_1 + \|P_m(\phi(P_n \mathbf{u}, P_n(\mathbf{u} * g' + \mathbf{d}')) - g')\|_1 \\ = & \|P_m(\phi(P_n \mathbf{u}, P_n \mathbf{y}) - g)\|_1 + \|P_m(\phi(P_n \mathbf{u}, P_n \mathbf{y}) - g')\|_1 \end{aligned} \quad (3.11)$$

$$\geq \|P_m(g - g')\|_1 \quad (3.12)$$

$$> D(\mathbf{u}, \mathcal{M}, \delta) - 2\epsilon \quad \text{by (3.10)} \quad (3.13)$$

Fix any $m > M$ and taking limits over n , we obtain

$$\begin{aligned} \limsup_{n \rightarrow \infty} \|P_m(\phi(P_n \mathbf{u}, P_n(\mathbf{u} * g + \mathbf{d})) - g)\|_1 & + \limsup_{n \rightarrow \infty} \|P_m(\phi(P_n \mathbf{u}, P_n(\mathbf{u} * g' + \mathbf{d}')) - g')\|_1 \\ & > D(\mathbf{u}, \mathcal{M}, \delta) - 2\epsilon \end{aligned} \quad (3.14)$$

Hence either

$$\limsup_{n \rightarrow \infty} \|P_m(\phi(P_n \mathbf{u}, P_n(\mathbf{u} * g + d)) - g)\|_1 > \frac{1}{2}D(\mathbf{u}, \mathcal{M}, \delta) - \epsilon \quad (3.15)$$

or

$$\limsup_{n \rightarrow \infty} \|P_m(\phi(P_n \mathbf{u}, P_n(\mathbf{u} * g' + d')) - g')\|_1 > \frac{1}{2}D(\mathbf{u}, \mathcal{M}, \delta) - \epsilon \quad (3.16)$$

Since this is true for every $\epsilon > 0$, we can conclude that $e_\infty^t(\phi, \mathcal{M}, \mathbf{u}, \delta) \geq \frac{1}{2}D(\mathbf{u}, \mathcal{M}, \delta)$. \square

The key question now is whether there exists an optimal algorithm which can always generate estimates with error converging to this lower bound. By the definition of the infinite-horizon uncertainty set, there exist two plants at a separation of $D(\mathbf{u}, \mathcal{M}, \delta)$ which can give rise to exactly the same output measurements. Thus in the worst case, there is no way for any finite-duration experiments to distinguish between them, and this gives rise to the lower bound proved above. Conversely, any two plants with a separation greater than $D(\mathbf{u}, \mathcal{M}, \delta)$ can be distinguished if we perform experiments of sufficiently long length. That is, if h is the true plant, and h' is another plant which is far away from h (separation greater than $D(\mathbf{u}, \mathcal{M}, \delta)$), there exists a time $T(h')$ for which one needs to observe the output to eliminate h' from consideration as a possible candidate. However, to guarantee that an accurate estimate at time n can be obtained, one needs $T(h') \leq n$ for *all* plants h' that are far away from h . Otherwise, although the identification algorithm always picks estimates which are consistent with the output seen so far, the estimates may nevertheless diverge from the true plant.

The issue discussed above is really one of *consistency* between finite-horizon experiments, where only finite data record is available for computing estimates, and infinite-horizon experiments, where the entire infinite data record is available. The question is when the latter can be viewed as a limit of the former. In [11], such a consistency result is established by placing a stationarity assumption on the noise and then appealing to the law of large numbers. As far as we know, this issue has not been considered in a unknown-but-bounded noise setting; in fact, it has been taken for granted that consistency always holds [9]. Instead, it will now be shown that some compactness condition on the model set is needed to guarantee consistency.

Definition 3.4 *A set Y is said to be σ -compact in a given topology if $Y = \cup_i Y_i$ where $Y_1 \subset Y_2 \subset Y_3 \dots$ are nested compact subsets of Y .*

The following theorem shows that, under a σ -compactness assumption on \mathcal{M} , $D(\mathbf{u}, \mathcal{M}, \delta)$ is an upper bound for the optimal asymptotic l_1 error. Combining with Proposition 3.3, we have upper and lower bounds that agree, within a factor of 2. Thus, the study of the optimal asymptotic l_1 error is reduced to the study of $D(\mathbf{u}, \mathcal{M}, \delta)$.

Theorem 3.5 *If the model set \mathcal{M} is σ -compact in the l_1 topology, then there is an identification algorithm ϕ^* such that $e_\infty^1(\phi^*, \mathcal{M}, \mathbf{u}, \delta) \leq D(\mathbf{u}, \mathcal{M}, \delta)$ for all \mathbf{u} and $\delta \geq 0$.*

Before proving Theorem 3.5, we need one more definition and a few lemmas.

Definition 3.6 *For given inputs \mathbf{u} and bound δ on disturbances, and $g \in \mathcal{X}$, define $T_{\mathbf{u}, \delta}(g)$ to be the smallest integer k such that $\|P_k(\mathbf{u} * g)\|_\infty > 2\delta$. If no such k exists, then $T_{\mathbf{u}, \delta}(g)$ is infinite.*

Lemma 3.7 *For any two plants $g, h \in \mathcal{M}$, $T_{\mathbf{u}, \delta}(g - h)$ is the smallest k such that there is no output \mathbf{y} with g and h in the same uncertainty set $S_k(\mathcal{M}, \mathbf{u}, \mathbf{y}, \delta)$.*

Proof. If $n = T_{\mathbf{u}, \delta}(g - h)$, then $\|P_n\{\mathbf{u} * (g - h)\}\|_\infty > 2\delta$, so for every output sequence \mathbf{y} , either $\|P_n\{\mathbf{u} * g - \mathbf{y}\}\|_\infty > \delta$ or $\|P_n\{\mathbf{u} * h - \mathbf{y}\}\|_\infty > \delta$, by the triangle inequality. Hence, g and h cannot be in the same uncertainty set $S_n(\mathcal{M}, \mathbf{u}, \mathbf{y}, \delta)$ for any \mathbf{y} . Conversely, if $n < T_{\mathbf{u}, \delta}(g - h)$, then $\|P_n\{\mathbf{u} * (g - h)\}\|_\infty \leq 2\delta$, so picking $\mathbf{y} = \mathbf{u} * (g + h)/2$ yields $\|P_n\{\mathbf{u} * g - \mathbf{y}\}\|_\infty \leq \delta$ and $\|P_n\{\mathbf{u} * h - \mathbf{y}\}\|_\infty \leq \delta$. Hence $g, h \in S_n(\mathcal{M}, \mathbf{u}, \mathbf{y}, \delta)$. \square

Thus, given two plants g and h , $T_{\mathbf{u}, \delta}(g - h)$ is the minimum duration for which one has to observe the output to ensure that at least one of the two plants can be eliminated from consideration as the true plant.

Lemma 3.8 *Let $g, h \in \mathcal{M}$. If $\|g - h\|_1 > D(\mathbf{u}, \mathcal{M}, \delta)$, then $T_{\mathbf{u}, \delta}(g - h) < \infty$.*

Proof. Suppose $T_{\mathbf{u}, \delta}(g - h) = \infty$. Then $\|P_k\{\mathbf{u} * (g - h)\}\|_\infty \leq 2\delta$ for every k , so $\|\mathbf{u} * (g - h)\|_\infty \leq 2\delta$. Now consider the disturbance $\mathbf{d} = \mathbf{u} * (h - g)/2$, and the infinite-horizon uncertainty set $S_\infty(\mathcal{M}, \mathbf{u}, \mathbf{u} * g + \mathbf{d}, \delta)$, arising when g is the true plant. (Note that $\|\mathbf{d}\|_\infty \leq \delta$.) But $\|\mathbf{u} * h - (\mathbf{u} * g + \mathbf{d})\|_\infty = \|\mathbf{u} * (h - g)/2\|_\infty \leq \delta$, so the plant h is also in the set $S_\infty(\mathcal{M}, \mathbf{u}, \mathbf{u} * g + \mathbf{d}, \delta)$. Hence, by definition of the worst-case infinite-horizon diameter, $\|g - h\|_1 \leq D(\mathbf{u}, \mathcal{M}, \delta)$. \square

Lemma 3.9 *Fix the inputs $\mathbf{u} \in \overline{Bl}_\infty^N$ and $\delta > 0$. Let $A \subset \mathcal{X}$ be compact in the l_1 topology, and suppose $T_{\mathbf{u}, \delta}(g)$ is finite for every $g \in A$. Then $\sup_{g \in A} T_{\mathbf{u}, \delta}(g)$ is also finite.*

Proof. Suppose $\sup_{g \in A} T_{\mathbf{u}, \delta}(g) = \infty$. Then there exists a sequence of plants $g^{(i)}$ in A such that $\lim_{i \rightarrow \infty} T_{\mathbf{u}, \delta}(g^{(i)}) = \infty$; furthermore, the sequence can be assumed to converge (in the l_1 topology) to a plant $g^* \in A$ since A is compact. Let $n^* \equiv T_{\mathbf{u}, \delta}(g^*) < \infty$. By definition, $\|P_{n^*}(\mathbf{u} * g^*)\|_\infty > 2\delta$.

But $\|P_{n^*}(\mathbf{u} * g)\|_\infty$ is a continuous function of g in the l_1 -topology, since $\|\mathbf{u}\|_\infty \leq 1$. Hence there exists a l_1 ball B around g^* such that for every $g' \in B$, $\|P_{n^*}(\mathbf{u} * g')\|_\infty > 2\delta$, i.e. $T_{\mathbf{u},\delta}(g') \leq n^*$ for every $g' \in B$. But this contradicts the fact that $\lim_{i \rightarrow \infty} T_{\mathbf{u},\delta}(g^{(i)}) = \infty$ since $g^{(i)} \rightarrow g^*$. Hence it can be concluded that $\sup_{g \in \mathcal{A}} T_{\mathbf{u},\delta}(g)$ is in fact finite. \square

Now we are in a position to prove Theorem 3.5.

Proof. Write \mathcal{M} as a nested union of compact subsets \mathcal{M}_i . Define the identification algorithm ϕ^* as follows: at each time n , the algorithm generates as an estimate by picking any arbitrary plant $\hat{h}^{(n)}$ in the set $S_n \cap \mathcal{M}_k$, where S_n is the uncertainty set after observing the output data until time n , and k is the least integer i such that $S_n \cap \mathcal{M}_i$ is non-empty. We claim that this algorithm will have an asymptotic error of at most $D(\mathbf{u}, \mathcal{M}, \delta)$ for all inputs \mathbf{u} and $\delta > 0$.

Fix the unknown plant $h \in \mathcal{M}$ and let $\epsilon > 0$. Also let \mathcal{M}_h be the smallest of the compact subsets \mathcal{M}_i 's which contains h . Define the set

$$A(h, \epsilon) \equiv \{g \in \mathcal{M}_h : \|g - h\|_1 \geq D(\mathbf{u}, \mathcal{M}, \delta) + \epsilon\} \quad (3.17)$$

and the number

$$T(h, \epsilon) \equiv \sup_{g \in A(h, \epsilon)} T_{\mathbf{u},\delta}(g - h) \quad (3.18)$$

Since $A(h, \epsilon)$ is a closed subset of \mathcal{M}_h (with respect to the l_1 -topology), it is also compact. By Lemma 3.8, $T_{\mathbf{u},\delta}(g - h)$ is finite for all $g \in A(h, \epsilon)$. Hence, by Lemma 3.9, $T(h, \epsilon)$ is also finite.

Now consider the estimates $\hat{h}^{(n)}$ generated by the algorithm ϕ^* . Since $\hat{h}^{(n)}$ is picked from the least k such that $S_n \cap \mathcal{M}_k$ is non-empty, $\hat{h}^{(n)}$ is guaranteed to be in \mathcal{M}_h for all n . (This is because $S_n \cap \mathcal{M}_h$ is non-empty: it contains the true plant h .) Also $\hat{h}^{(n)}$ is in the uncertainty set S_n and by Lemma 3.7, $T_{\mathbf{u},\delta}(\hat{h}^{(n)} - h) > n$. If we now take any $n > T(h, \epsilon)$, we have $T_{\mathbf{u},\delta}(\hat{h}^{(n)} - h) > T(h, \epsilon)$ so $\hat{h}^{(n)}$ is not in $A(h, \epsilon)$. But $\hat{h}^{(n)}$ is in \mathcal{M}_h , so it follows that $\|\hat{h}^{(n)} - h\|_1 < D(\mathbf{u}, \mathcal{M}, \delta) + \epsilon$.

Since ϵ is arbitrary, it can now be concluded that

$$\limsup_{n \rightarrow \infty} \|\hat{h}^{(n)} - h\|_1 \leq D(\mathbf{u}, \mathcal{M}, \delta)$$

completing the proof. \square

The above construction of the asymptotically near-optimal algorithm ϕ^* can be viewed as an application of *Occam's Razor* – that one should always use the “simplest” theory to explain the given data. Here, as is true in general, there is no absolute measure of simplicity. Rather it is defined by the choice of the nested partitioning of the model set, $\mathcal{M} = \cup_i \mathcal{M}_i$. Given this nested structure,

plants in the smaller \mathcal{M}_i 's are considered to be simpler than those in larger \mathcal{M}_i . Convergence of the estimates is guaranteed by always choosing the simplest plant that is consistent with the data seen so far. This avoids overfitting of data, a problem which crops up all the time in statistics and pattern recognition. It is interesting to note that this same principle of Occam's Razor has also been applied to guarantee convergence in distribution-free probabilistic learning problems [1, 20].

A similar condition on the model set for convergence in the *truncated* l_1 seminorm can be derived. It turns out that a σ -compactness condition is also sufficient, but in a different topology on $\mathcal{X} \equiv \mathfrak{R}^\omega$, the so called product topology [18]. (Here, ω is the set of all natural numbers.)

Definition 3.10 *The product topology on \mathfrak{R}^ω has as its basis all sets of the form $\prod_{n=0}^{\infty} U_n$ where each U_n is an open interval in \mathfrak{R} and $U_n = \mathfrak{R}$ except for finitely many n .*

This topology is strictly weaker than the l_1 -topology on \mathfrak{R}^ω . The product topology is relevant because convergence in this topology is the same as component-wise convergence, which is essentially the notion of convergence captured by using the truncated l_1 seminorm.

We will need the following lemma, which is analogous to Lemma 3.9.

Lemma 3.11 *Fix the inputs $\mathbf{u} \in \overline{BI}_\infty^N$ and $\delta > 0$. Let $A \subset \mathcal{X}$ be compact in the product topology, and suppose $T_{\mathbf{u},\delta}(g)$ is finite for every $g \in A$. Then $\sup_{g \in A} T_{\mathbf{u},\delta}(g)$ is also finite.*

Proof. By examining the proof of Lemma 3.11, one sees that the lemma holds for the l_1 topology because $\|P_n(\mathbf{u} * g)\|_\infty$ is a continuous function of g in that topology. However, $\|P_n(\mathbf{u} * g)\|_\infty$ is also continuous in the *product* topology, since it depends only on the first n components of g . Hence the same lemma holds for the product topology. \square

The sufficient condition on \mathcal{M} for asymptotic convergence, in the truncated l_1 seminorm, to the worst-case infinite-horizon diameter can now be stated and proved.

Theorem 3.12 *If the model set \mathcal{M} is σ -compact in the product topology, then there is an identification algorithm ϕ^* such that $e_\infty^t(\phi^*, \mathcal{M}, \mathbf{u}, \delta) \leq D(\mathbf{u}, \mathcal{M}, \delta)$ for all \mathbf{u} and $\delta \geq 0$.*

Proof. Write \mathcal{M} as a nested union of subsets \mathcal{M}_i compact in the product topology. We claim that the algorithm ϕ^* introduced in the proof of Theorem 3.5 will also do here.

Fix the unknown plant $h \in \mathcal{M}$ and define \mathcal{M}_h to be the smallest of the compact subsets \mathcal{M}_i which contains h . Take any integer $m \geq 0$ and let $\epsilon > 0$. Define

$$A_m(h, \epsilon) \equiv \{g \in \mathcal{M}_h : \|P_m(g - h)\|_1 \geq D(\mathbf{u}, \mathcal{M}, \delta) + \epsilon\} \quad (3.19)$$

$$T_m(h, \epsilon) \equiv \sup_{g \in A_m(h, \epsilon)} T_{\mathbf{u}, \delta}(g - h) \quad (3.20)$$

Now the truncated l_1 ball $\{g \in \mathfrak{R}^\omega : \|P_m(g - h)\|_1 < D(\mathbf{u}, \mathcal{M}, \delta) + \epsilon\}$ is an open set in the product topology. Therefore $A_m(h, \epsilon)$ is a closed subset of \mathcal{M}_h and is also compact in the product topology. Moreover, by Lemma 3.8, $T_{\mathbf{u}, \delta}(g - h)$ is finite for all $g \in A_m(h, \epsilon)$. Hence, it follows from Lemma 3.11 that $T_m(h, \epsilon)$ is finite. Now, using exactly the same argument as in the proof of Theorem 3.5, it can be shown that, for every $n > T_m(h, \epsilon)$, the estimates $\hat{h}^{(n)}$ satisfy

$$\|P_m(\hat{h}^{(n)} - h)\|_1 < D(\mathbf{u}, \mathcal{M}, \delta) + \epsilon \quad (3.21)$$

Since ϵ is arbitrary, it can be concluded that

$$\limsup_{n \rightarrow \infty} \|P_m(\hat{h}^{(n)} - h)\|_1 \leq D(\mathbf{u}, \mathcal{M}, \delta) \quad (3.22)$$

and since this is true for all m , $e_\infty^t(\phi^*, \mathcal{M}, \mathbf{u}, \delta) \leq D(\mathbf{u}, \mathcal{M}, \delta)$. \square

It is interesting to compare the topological conditions on the model set \mathcal{M} for convergence in the l_1 norm and for convergence in the truncated l_1 seminorm. By Proposition 2.3, the asymptotic truncated l_1 error is always less than the asymptotic l_1 error. It therefore follows that the σ -compactness of \mathcal{M} in the l_1 -topology, which is sufficient for convergence in the l_1 norm, is also a sufficient condition for convergence in the truncated l_1 norm. However, most model sets containing unstable plants are not σ -compact in the l_1 -topology and therefore the condition is too strong to be useful. On the other hand, it is much easier to satisfy the condition of σ -compactness in the *product* topology, and hence it is more useful for analysing model sets containing unstable plants. This fact will become evident when we later apply these general results on specific model sets.

The main results of this section are summarized in the following theorem.

Theorem 3.13 *If the model set \mathcal{M} is σ -compact in the l_1 topology, then the optimal asymptotic truncated l_1 error and the worst-case infinite-horizon diameter are related by*

$$\frac{1}{2}D(\mathbf{u}, \mathcal{M}, \delta) \leq E_\infty^1(\mathbf{u}, \mathcal{M}, \delta) \leq D(\mathbf{u}, \mathcal{M}, \delta)$$

for all inputs \mathbf{u} and $\delta \geq 0$. If the model set \mathcal{M} is σ -compact in the product topology, then the optimal asymptotic truncated l_1 error and the worst-case infinite-horizon diameter are related by

$$\frac{1}{2}D(\mathbf{u}, \mathcal{M}, \delta) \leq E_\infty^t(\mathbf{u}, \mathcal{M}, \delta) \leq D(\mathbf{u}, \mathcal{M}, \delta)$$

4 Analysis of Infinite-Horizon Uncertainty

Theorem 3.13 states that the optimal worst-case asymptotic error achievable by any identification algorithm is characterized by the function $D(\mathbf{u}, \mathcal{M}, \delta)$, measuring the worst-case uncertainty from infinite-horizon experiments. It describes the intrinsic difficulty of identifying plants in a given model set, independent of the specific identification algorithm used. This result enables us to move from the analysis of the error of specific algorithms to the analysis of the function $D(\mathbf{u}, \mathcal{M}, \delta)$. In particular, we would like to determine conditions on the model set \mathcal{M} for the existence of inputs \mathbf{u} such that $D(\mathbf{u}, \mathcal{M}, \delta)$ is finite, or even better, $D(\mathbf{u}, \mathcal{M}, \delta)$ vary continuously with the noise bound δ at $\delta = 0$. This would imply that identification accuracy is *robust* to measurement noise.

We shall first prove a general lower bound for $D(\mathbf{u}, \mathcal{M}, \delta)$ which holds for all inputs \mathbf{u} and for any non-trivial model set.

Proposition 4.1 *Assume that there exist $g, h \in \mathcal{M}$ such that $\|g - h\|_1 = 2\delta$. Then*

$$D(\mathbf{u}, \mathcal{M}, \delta) \geq 2\delta \quad \forall N > 0, \forall \mathbf{u} \in \overline{BI}_\infty^N$$

Proof. Let $g, h \in \mathcal{M}$ satisfy $\|g - h\|_1 = 2\delta$. Suppose that \mathbf{u} are the inputs used in the identification experiments and h is the actual plant. Let the disturbance be $\mathbf{d} = \mathbf{u} * (g - h)/2$. Note that $\|\mathbf{d}\|_\infty \leq \|\mathbf{u}\|_\infty \|\frac{g-h}{2}\|_1 = \delta$.

The observed output is $\mathbf{y} = \mathbf{u} * h + \mathbf{d} = \mathbf{u} * (g + h)/2$. Now, $\|\mathbf{u} * g - \mathbf{y}\|_\infty = \|\frac{1}{2}\mathbf{u} * (g - h)\|_\infty \leq \frac{1}{2}\|\mathbf{u}\|_\infty \|g - h\|_1 \leq \delta$. Therefore, $g \in S_\infty(\mathcal{M}, \mathbf{u}, \mathbf{y}, \delta)$. Since h is also in $S_\infty(\mathcal{M}, \mathbf{u}, \mathbf{y}, \delta)$, it follows that

$$\text{diam}(S_\infty(\mathcal{M}, \mathbf{u}, \mathbf{y}, \delta)) \geq \|g - h\|_1 = 2\delta$$

Since $D(\mathbf{u}, \mathcal{M}, \delta)$ is the diameter of the largest possible uncertainty set, the desired lower bound follows. \square

The value of the worst-case diameter $D(\mathbf{u}, \mathcal{M}, \delta)$ is in general difficult to evaluate because it is the supremum over the diameter of all possible infinite-horizon uncertainty sets. However, it turns out that for an important class of model sets, $D(\mathbf{u}, \mathcal{M}, \delta)$ has a simple characterization. These are the model sets which are convex and balanced. (A set A is said to be balanced if for every h in A , $-h$ is also in A .) The following proposition gives the characterization, and is a variation of a basic result in information-based complexity theory [14].

Proposition 4.2 *If \mathcal{M} is a balanced convex subset of \mathcal{X} , then the worst-case diameter is attained when the true plant and the disturbance are both 0. That is,*

$$D(\mathbf{u}, \mathcal{M}, \delta) \equiv \sup_{h \in \mathcal{M}} \sup_{\|\mathbf{d}\|_\infty \leq \delta} \text{diam}(S_\infty(\mathcal{M}, \mathbf{u}, \mathbf{u} * h + \mathbf{d}, \delta)) = \text{diam}(S_\infty(\mathcal{M}, \mathbf{u}, 0, \delta))$$

Proof. Consider an arbitrary uncertainty set $S_\infty(\mathcal{M}, \mathbf{u}, \mathbf{y}, \delta)$. Suppose $g', g'' \in S_\infty(\mathcal{M}, \mathbf{u}, \mathbf{y}, \delta)$. Then $\mathbf{y} = \mathbf{u} * g' + \mathbf{d}' = \mathbf{u} * g'' + \mathbf{d}''$ for some disturbances $\mathbf{d}', \mathbf{d}''$ satisfying $\|\mathbf{d}'\|_\infty \leq \delta, \|\mathbf{d}''\|_\infty \leq \delta$. Hence, $\|\mathbf{u} * (g' - g'')\|_\infty \leq 2\delta$. Therefore, the diameter of the uncertainty set satisfies

$$\begin{aligned} \text{diam}(S_\infty(\mathcal{M}, \mathbf{u}, \mathbf{y}, \delta)) &\equiv \sup_{g', g'' \in S_\infty(\mathcal{M}, \mathbf{u}, \mathbf{y}, \delta)} \|g' - g''\|_1 \\ &\leq \sup_{\|\mathbf{u} * (g' - g'')\|_\infty \leq 2\delta} \|g' - g''\|_1 \end{aligned} \quad (4.23)$$

Because \mathcal{M} is convex and balanced, $(g' - g'')/2 \in \mathcal{M}$. Noting that

$$S_\infty(\mathcal{M}, \mathbf{u}, 0, \delta) = \{g \in \mathcal{M} : \|\mathbf{u} * g\|_\infty \leq \delta\} \quad (4.24)$$

and writing $g = (g' - g'')/2$, it follows from above that

$$\text{diam}(S_\infty(\mathcal{M}, \mathbf{u}, \mathbf{d}, \delta)) \leq \sup_{g \in S_\infty(\mathcal{M}, \mathbf{u}, 0, \delta)} \|2g\|_1 = \text{diam}(S_\infty(\mathcal{M}, \mathbf{u}, 0, \delta))$$

where the last equality follows because $S_\infty(\mathcal{M}, \mathbf{u}, 0, \delta)$ is a balanced set. The desired result follows.

□

We shall now obtain necessary and sufficient conditions on balanced convex model sets for the existence of inputs \mathbf{u} such that $D(\mathbf{u}, \mathcal{M}, \delta)$ is finite. These conditions are given in terms of the notion of *stability testing*, which we will now define.

Recall that a plant h is BIBO stable if for every input u in l_∞ , the output y is also in l_∞ ; or, equivalently, if h is in l_1 . Note that stability is defined in terms of the behaviour of a plant on an *infinite* class of inputs. It is of interest to inquire whether for some model sets $\mathcal{M} \subset \mathcal{X}$, the stability of plants in \mathcal{M} can be determined in terms of their behaviour on only a finite set of inputs. This motivates the following definition.

Definition 4.3 *A model set \mathcal{M} is said to be finitely testable for stability if there exists some $N > 0$ and inputs $u^{(1)}, u^{(2)}, \dots, u^{(N)} \in \overline{Bl}_\infty$ such that $h \in \mathcal{M}$ is BIBO stable if and only if $u^{(i)} * h \in l_\infty$ for all $i = 1, 2, \dots, N$. Equivalently, in terms of the vectorial notation, there exist $N > 0$ and $\mathbf{u} \in \overline{Bl}_\infty^N$ such that $h \in \mathcal{M}$ is stable if and only if $\mathbf{u} * h \in \overline{Bl}_\infty^N$.*

Theorem 4.4 *Assume \mathcal{M} is balanced and convex. Then there exist inputs \mathbf{u} such that $D(\mathbf{u}, \mathcal{M}, \delta)$ is finite if and only if it is finitely testable for stability. Moreover, if N inputs are sufficient for testing stability, then there exist $N + 1$ inputs \mathbf{u} such that*

$$D(\mathbf{u}, \mathcal{M}, \delta) \leq 2\delta$$

Proof. By Proposition 4.2,

$$D(\mathbf{u}, \mathcal{M}, \delta) = \text{diam}(S_\infty(\mathcal{M}, \mathbf{u}, 0, \delta)) \equiv \sup_{g \in \mathcal{M}, \|\mathbf{u} * g\|_\infty \leq \delta} 2\|g\|_1$$

Suppose that \mathcal{M} is not finitely testable for stability. Then, given any $N > 0$ and any $\mathbf{u} \in \overline{Bl}_\infty^N$, there exists $g \in \mathcal{M}$ such that g is unstable but $\mathbf{u} * g \in l_\infty^N$. Define

$$g' = \min\left\{1, \frac{\delta}{\|\mathbf{u} * g\|_\infty}\right\}g \quad (4.25)$$

We have $g' \in \mathcal{M}$ since \mathcal{M} is balanced and convex. Furthermore,

$$\|\mathbf{u} * g'\|_\infty \leq \frac{\delta}{\|\mathbf{u} * g\|_\infty} \|\mathbf{u} * g\|_\infty = \delta. \quad (4.26)$$

Hence $g' \in S_\infty(\mathcal{M}, \mathbf{u}, 0, \delta)$. But g' is unstable, so $\|g'\|_1 = \infty$. Therefore, $\text{diam}(S_\infty(\mathcal{M}, \mathbf{u}, 0, \delta)) = \infty$ and $D(\mathbf{u}, \mathcal{M}, \delta) = \infty$ for all inputs \mathbf{u} and $\delta > 0$.

Conversely, suppose that \mathcal{M} is finitely testable for stability. Then, by definition, there exist $N > 0$ and $\mathbf{u} = [u^{(1)}, \dots, u^{(N)}] \in \overline{Bl}_\infty^N$ such that

$$\mathbf{u} * g \in l_\infty^N \Leftrightarrow g \in l_1 \quad \forall g \in \mathcal{M} \quad (4.27)$$

Hence, $S_\infty(\mathcal{M}, \mathbf{u}, 0, \delta) \subset l_1$ for every $\delta \geq 0$.

Now, let \mathcal{A} be the set of all *finite* sequences of 1's and -1's:

$$\mathcal{A} \equiv \{(a_1, a_2, \dots, a_k) : k \geq 1, a_i \in \{1, -1\}, \forall i\} \quad (4.28)$$

Let $v \in \overline{Bl}_\infty$ be an input sequence such that for every finite sequence $a \in \mathcal{A}$, there exist m, n such that

$$(v_m, v_{m+1}, \dots, v_{m+n}) = a \quad (4.29)$$

The sequence v is said to *contain* all finite sequences of 1's and -1's. Clearly, since \mathcal{A} is only countable, there exists a v with such a property. (For example, one can construct such a v by simply concatenating the elements of \mathcal{A} .)

It will now be shown that performing $N + 1$ experiments using the inputs \mathbf{u} (which can test stability) and the additional input v , the worst case diameter can be made finite.

Let $\mathbf{u}' = [u^{(1)}, \dots, u^{(N)}, v]$ be the augmented input sequence vector. Performing an additional experiment with input v has the effect of reducing the zero-output uncertainty set:

$$S_\infty(\mathcal{M}, \mathbf{u}', 0, \delta) = S_\infty(\mathcal{M}, \mathbf{u}, 0, \delta) \cap S_\infty(\mathcal{M}, v, 0, \delta) \quad (4.30)$$

Consider any $g \in S_\infty(\mathcal{M}, \mathbf{u}', 0, \delta)$. We see that g is also in $S_{\text{infty}}(\mathcal{M}, \mathbf{u}, 0, \delta)$ and so $\|g\|_1 < \infty$. Hence for any $\epsilon > 0$ there exists M such that

$$\sum_{k=M+1}^{\infty} |g_k| < \epsilon \quad (4.31)$$

Now consider the finite sequence

$$(\text{sgn}(g_M), \text{sgn}(g_{M-1}), \dots, \text{sgn}(g_0)) \in \mathcal{A}$$

where sgn is the signum function such that $\text{sgn}(x) = 1$ if $x \geq 0$ and $\text{sgn}(x) = -1$ if $x < 0$.

By definition of the sequence v , there exists m such that

$$v_m = \text{sgn}(g_M), v_{m+1} = \text{sgn}(g_{M-1}), \dots, v_{m+M} = \text{sgn}(g_0)$$

We then have

$$\begin{aligned} |(v * g)_{m+M}| &= \left| \sum_{k=0}^{m+M} v_{m+M-k} g_k \right| \\ &= \left| \sum_{k=0}^M v_{m+M-k} g_k + \sum_{k=M+1}^{m+M} v_{m+M-k} g_k \right| \\ &= \left| \sum_{k=0}^M \text{sgn}(g_k) g_k + \sum_{k=M+1}^{m+M} v_{m+M-k} g_k \right| \\ &\geq \sum_{k=0}^M |g_k| - \sum_{k=M+1}^{m+M} |g_k| \\ &\geq \|g\|_1 - 2\epsilon, \end{aligned} \quad (4.32)$$

But $g \in S_\infty(\mathcal{M}, v, 0, \delta)$, so $|(v * g)_{m+M}| \leq \delta$. Hence it follows from inequality (4.32) that $\|g\|_1 \leq \delta + 2\epsilon$. Since this is true for every $\epsilon > 0$, it follows that $\|g\|_1 \leq \delta$ for any $g \in S_\infty(\mathcal{M}, \mathbf{u}', 0, \delta)$. Thus,

$$D(\mathbf{u}, \mathcal{M}, \delta) = \text{diam}(S_\infty(\mathcal{M}, \mathbf{u}', 0, \delta)) = \sup_{g \in S_\infty(\mathcal{M}, \mathbf{u}', 0, \delta)} 2\|g\|_1 \leq 2\delta$$

□

Theorem 4.4 gives an interesting dichotomy of the class of all balanced, convex model sets. It shows that for a balanced convex model set \mathcal{M} , either $D(\mathbf{u}, \mathcal{M}, \delta) = \infty$ for every \mathbf{u} and for every $\delta > 0$, or $D(\mathbf{u}, \mathcal{M}, \delta) \leq 2\delta$ for some \mathbf{u} . In view of the general lower bound of 2δ on $D(\mathbf{u}, \mathcal{M}, \delta)$ for most model sets (Proposition 4.1), this implies that balanced convex model sets can either be identified very accurately or with infinite error regardless of how much experimentation one does.

It is also of interest to note that the random binary sequence, a commonly used identification input generated by randomly and independently picking each value to be 1 or -1 , has the desired property of containing all finite sequences of 1's and -1 's, with probability 1.

5 Analysis of Specific Model Sets

We shall now apply the general results obtained in the previous two sections to analyse the optimal asymptotic errors for specific model sets. First, we present a general lower bound on the optimal asymptotic errors that follows directly from Theorem 3.13 and Proposition 4.1.

Proposition 5.1 *Let $\delta \geq 0$. Assume that there exists $g, h \in \mathcal{M}$ such that $\|g - h\|_1 = 2\delta$. Then*

$$E_{\infty}^1(\mathbf{u}, \mathcal{M}, \delta) \geq \delta \quad \text{and} \quad E_{\infty}^t(\mathbf{u}, \mathcal{M}, \delta) \geq \delta \quad \forall N > 0, \mathbf{u} \in \overline{Bl}_{\infty}^N$$

Thus δ is an intrinsic lower bound to the identification accuracy for most model sets. Identification can be considered successful if the optimal worst-case asymptotic error turns out to be close to δ .

Next we will look at three specific model sets: the set of all stable plants, the set of all finite-dimensional (stable and unstable) plants, and the set \mathcal{X} of all LTI systems.

A. Stable Plants

Let \mathcal{M}_{stab} be the set of all BIBO stable plants. (The set \mathcal{M}_{stab} can be identified with the space l_1). To obtain tight bounds on the optimal asymptotic error, we begin by investigating its topological properties.

Lemma 5.2 *Let g be a stable plant. The set of uniformly stable plants:*

$$\mathcal{M}_s(g) \equiv \{h : |h_k| \leq |g_k| \quad \forall k\}$$

is compact in the l_1 topology.

Proof. Since $\mathcal{M}_s(g)$ is a closed subset of l_1 , it is complete. Therefore, to prove compactness, it suffices to prove that the set is totally bounded, that is, $\mathcal{M}_s(g)$ can be covered by a finite number of l_1 balls of radius ϵ for every $\epsilon > 0$.

Let $\epsilon > 0$. Pick N such that $\sum_{k=N+1}^{\infty} |g_k| < \epsilon/2$. For each k between 0 and N , partition the interval $[-|g_k|, |g_k|]$ by choosing points $x_{k0} < x_{k1} < \dots < x_{kM}$ such that $x_{k0} = -|g_k|$ and $x_{kM} = |g_k|$ and so that adjacent points differ by no more than ϵ/N . (Here, M is some appropriately chosen number). Consider the finite set of plants in \mathcal{M} defined by

$$F \equiv \{h : h_k \in \{x_{k0}, x_{k1}, \dots, x_{kM}\} \text{ for } k \leq N \text{ and } h_k = 0 \text{ for } k > N\} \quad (5.33)$$

Given any $h \in \mathcal{M}_s(g)$, there exists some $h' \in F$ such that $|h'_k - h_k| \leq \epsilon/2N$ for all $k \leq N$. Hence

$$\begin{aligned} \|h' - h\|_1 &= \sum_{k=0}^N |h'_k - h_k| + \sum_{k=N+1}^{\infty} |h'_k - h_k| \\ &\leq \sum_{k=0}^N \epsilon/2N + \epsilon/2 = \epsilon \end{aligned}$$

Hence the finite number of balls of radius ϵ centred at elements of F cover $\mathcal{M}_s(g)$. \square

The following is the main result about the set of all stable systems.

Proposition 5.3 *There is a single experiment, using an input $u^* \in \overline{Bl}_\infty$, such that for every $\delta \geq 0$, the optimal asymptotic l_1 error satisfies*

$$E_\infty^1(u^*, \mathcal{M}_{stab}, \delta) \leq 2\delta$$

Proof. For each $i \geq 1$, define a plant $g^{(i)}$ such that $g_k = i(1-1/i)^k$ for every k . Let $\mathcal{M}_i \equiv \mathcal{M}_s(g^{(i)})$ and $\mathcal{M}_\infty \equiv \cup_i \mathcal{M}_i$. It is easy to see that the \mathcal{M}_i 's are nested and that \mathcal{M}_{stab} is the l_1 closure of \mathcal{M}_∞ . Since each of the \mathcal{M}_i 's is compact by the previous lemma, \mathcal{M}_∞ is σ -compact and by Theorem 3.5, there is an algorithm ϕ^* such that for every $\delta \geq 0$ and inputs \mathbf{u} ,

$$e_\infty^1(\phi^*, \mathcal{M}_\infty, \mathbf{u}, \delta) \leq D(\mathbf{u}, \mathcal{M}_\infty, \delta) \quad (5.34)$$

Now,

$$\begin{aligned} e_\infty^1(\phi^*, \mathcal{M}_{stab}, \mathbf{u}, \delta) &\leq \lim_{x \downarrow \delta} e_\infty^1(\phi^*, \mathcal{M}_\infty, \mathbf{u}, x) \quad \text{by Prop. 2.4} \\ &\leq \lim_{x \downarrow \delta} D(\mathbf{u}, \mathcal{M}_\infty, x) \end{aligned} \quad (5.35)$$

Now, \mathcal{M}_{stab} is balanced and convex. Since it contains only stable plants, clearly no input is required to test stability. By Theorem 4.4, then, any single input u^* which contains all possible finite sequences of 1's and -1 's will yield $D(u^*, \mathcal{M}_{stab}, x) \leq 2x$ for all $x \geq 0$. Thus

$$e_{\infty}^1(\phi^*, \mathcal{M}_{stab}, u^*, \delta) \leq 2\delta$$

□

Hence, to identify a plant accurately in the limit, it is enough to know *a priori* that it is stable; no additional information, such as bounds on decay rate and gain, is necessary. The achievable accuracy varies continuously with the noise bound δ for small δ ; thus, identification can be performed robust to measurement noise.

B. Finite-Dimensional Plants

We will now consider the model set \mathcal{M}_{fd} , the set of all stable and unstable plants with rational z -transform. Since it contains unstable systems, it is unreasonable to discuss convergence of error in the l_1 norm. Instead, we will concentrate on analysing the asymptotic *truncated* l_1 error.

First we show that \mathcal{M}_{fd} is σ -compact in the product topology.

Proposition 5.4 *Let $p, q, K \geq 0$ and $\mathcal{M}_{fd}(p, q, K)$ be the class of all finite-dimensional systems having z -transforms⁴*

$$\frac{b_p z^p + b_{p-1} z^{p-1} + \dots + b_0}{z^q + a_{q-1} z^{q-1} + \dots + a_0}$$

with bounded parameters: $|a_i| \leq K$ and $|b_i| \leq K$ for all i . $\mathcal{M}_{fd}(p, q, K)$ is compact in the product topology.

Proof. If we define $\theta = [b_p, \dots, b_0, a_{q-1}, \dots, a_0]^t$ and $\mathbf{z}_n = [u_{n-p}, \dots, u_n, y_{n-q+1}, \dots, y_{n-1}]$, then the input u and output y of an arbitrary plant in $\mathcal{M}_{fd}(p, q, K)$ are related via the difference equation

$$y_n = \theta^t \mathbf{z}_n \tag{5.36}$$

Since the parameters are bounded, θ lies in a compact subset $A \subset \mathfrak{R}^{p+q+1}$. Let $f : A \rightarrow \mathcal{M}_{fd}(p, q, K)$ be the parameterization of plants in $\mathcal{M}_{fd}(p, q, K)$ via Equation (5.36), and $f_i(\theta)$ be the i th component of the impulse response of the plant $f(\theta)$. It can be deduced from Equation (5.36) that each f_i is a polynomial in θ and hence is continuous. Therefore the parameterization

⁴In this paper, the z -transform of a system with impulse response h is $\sum_{i=0}^{\infty} h_i z^i$.

f is continuous with respect to the product topology on $\mathcal{M}_{fd}(p, q, K)$. Since the domain of f is compact, its image $\mathcal{M}_{fd}(p, q, K)$ is compact in the product topology. \square

Proposition 5.5 \mathcal{M}_{fd} is σ -compact in the product topology.

Proof. For each $i \geq 1$, define $\mathcal{M}_i = \cup_{p+q \leq i} \mathcal{M}_{fd}(p, q, i)$. Each \mathcal{M}_i is compact in the product topology since it is a finite union of compact sets. Also, it is clear that $\mathcal{M}_{fd} = \cup_i \mathcal{M}_i$, and hence it is σ -compact. \square

It follows that Theorem 3.13 can be applied on \mathcal{M}_{fd} and the worst-case infinite-horizon diameter $D(\mathbf{u}, \mathcal{M}_{fd}, \delta)$ characterizes the optimal asymptotic error $E_\infty^t(\mathbf{u}, \mathcal{M}_{fd}, \delta)$. Moreover, since \mathcal{M}_{fd} is balanced and convex, it suffices, by virtue of Theorem 4.4, to verify if \mathcal{M}_{fd} is finitely testable for stability. We begin with two definitions.

Definition 5.6 For a sequence $u \in l_\infty$, let $Z(u)$ denote the set of all zeros of its z -transform $U(z)$ inside the open unit disk. (Note that $U(z)$ is analytic inside the open unit disk.)

Definition 5.7 A sequence u is said to excite at frequency $\omega \in [0, 2\pi]$ if

$$\limsup_{n \rightarrow \infty} \left| \sum_{k=0}^n u_k e^{-jk\omega} \right| = \infty$$

i.e. the Fourier series of u at ω is unbounded. Let $\Omega(u)$ denote the set of all frequencies at which u excites.

We shall now give the following key result

Theorem 5.8 \mathcal{M}_{fd} is testable for stability by bounded inputs $u^{(1)}, \dots, u^{(N)}$ if and only if the inputs have the following properties:

- 1) $\bigcup_{i=1}^N \Omega(u^{(i)}) = [0, 2\pi]$
- 2) $\bigcap_{i=1}^N Z(u^{(i)}) = \emptyset$

Hence the inputs can test the stability of finite-dimensional plants if they excite at all frequencies and have no common zeros in the unit disk.

To prove this result, we need the following lemma, the proof of which is elementary but tedious, and will be omitted.

Lemma 5.9 *Let $u \in \overline{B}l_\infty$ and let h be a complex-valued impulse response (i.e. the sequence values can be complex) with a strictly proper rational transfer function*

$$H(z) = \frac{\sum_{i=0}^{M-1} \alpha_i z^i}{(z - e^{j\omega})^M}$$

(It has a single pole repeated M times at $e^{j\omega}$). Then:

- 1) *If u excites at frequency ω , the output $u * h$ is unbounded.*
- 2) *If u does not excite at ω and $M = 1$ (the pole is simple), the output $u * h$ is bounded.*

Armed with this lemma, we can now prove Theorem 5.8.

Proof. (if part)

Let $u^{(1)}, u^{(2)}, \dots, u^{(N)} \in \overline{B}l_\infty$ be N inputs satisfying properties (1) and (2). Let $h \in \mathcal{M}_{fd}$ with a rational z -transform $H(z)$, and assume that the outputs $u^{(i)} * h, i = 1, \dots, N$, are all bounded. We shall show that h must be stable.

Suppose that $H(z)$ has a pole $z = z_1$ inside the open unit disk. Since the inputs have no common zeros, then one of the inputs, say $u^{(i)}$, has no zero at $z = z_1$. Hence the output $y^{(i)}$ must have a pole at $z = z_1$, and therefore cannot be bounded.

Thus, H can only have poles on or outside of the unit circle. Write

$$H(z) = H_u(z) + H_s(z) \tag{5.37}$$

where $H_s(z)$ contains the stable poles (outside the unit circle) and the finite impulse response (FIR) part of $H(z)$, and $H_u(z)$ is strictly proper with all poles on the unit circle. Let h_u and h_s be the inverse transforms of H_u and H_s , respectively. Since the output $u * h_s$ corresponding to the stable part must be bounded, one needs only to verify that the boundedness of $u^{(i)} * h_u$ for every i implies $h_u = 0$.

Suppose that H_u is not identically 0 and has $L > 0$ poles (counting multiplicities) on the unit circle at distinct frequencies $\omega_1, \omega_2, \dots, \omega_M$. Then $H_u(z)$ can be decomposed as

$$H_u(z) = \sum_{i=1}^M H_i(z) \tag{5.38}$$

where

$$H_i(z) = \frac{\sum_{k=0}^{L_i-1} \alpha_{ik} z^k}{(z - e^{j\omega_i})^{L_i}} \tag{5.39}$$

and L_i is the order of the pole at $z = e^{j\omega_i}$.

Consider a minimal state space realization of the system with transfer function $H_u(z)$, where the states \mathbf{x} consist of the modes corresponding to each pole of the system. The dimension of the realization is L and some of the states are complex but they occur in conjugate pairs. (These correspond to conjugate poles.) Since $\cup_{i=1}^N \Omega(u^{(i)}) = [0, 2\pi]$, the frequency ω_1 lies in $\Omega(v)$ for some input $v \in \{u^{(1)}, \dots, u^{(N)}\}$. By Proposition 5.9,

$$\mathbf{y}^{(1)} = v * h^{(1)} \notin l_\infty \quad (5.40)$$

where $h^{(1)}$ is the impulse response whose z-transform is $H_1(z)$.

If $\mathbf{x}^{(1)}$ are the modal states (of dimension L_1) corresponding to this pole at ω_1 , the system $h^{(1)}$ can be realized minimally as

$$\mathbf{x}_{n+1}^{(1)} = A_1 \mathbf{x}_n^{(1)} + B_1 v_n, \quad y_n^{(1)} = C_1 \mathbf{x}_n^{(1)} \quad (5.41)$$

for some matrices A_1, B_1, C_1 .

Since $y^{(1)}$ is unbounded but v is bounded, it follows from Equation 5.41 that the modal states $\mathbf{x}^{(1)}$ must be unbounded given input v . But the overall state \mathbf{x} for the entire system $H_u(z)$ is an aggregation of the modal states and hence must become unbounded too when input v is applied. The last step is to show that this implies that the output of the overall system must be unbounded also.

Let the minimal state space realization of H_u be

$$\mathbf{x}_{n+1} = A \mathbf{x}_n + B v_n, \quad y_n = C \mathbf{x}_n \quad (5.42)$$

From Equation 5.42, a sequence of equations is obtained as

$$\begin{aligned} y_n &= C \mathbf{x}_n \\ y_{n+1} &= C A \mathbf{x}_n + C B v_n \\ &\vdots \\ y_{n+L-1} &= C A^{L-1} \mathbf{x}_n + \sum_{i=0}^{L-2} C A^i B v_n \end{aligned}$$

Let

$$\mathbf{y}_n = \begin{bmatrix} y_n \\ y_{n+1} \\ \vdots \\ y_{n+L-1} \end{bmatrix}, Q_0(A, C) = \begin{bmatrix} C \\ C A \\ \vdots \\ C A^{L-1} \end{bmatrix}, E = \begin{bmatrix} 0 \\ C B \\ \vdots \\ \sum_{i=0}^{L-2} C A^i B \end{bmatrix}$$

The sequence of output equations can then be written as

$$\mathbf{y}_n = Q_0(A, C) \mathbf{x}_n + E v_n \quad (5.43)$$

Note that $Q_0(A, C)$ is the observability matrix of the system. By the minimality of the realization, $Q_0(A, C)$ is invertible. Since \mathbf{x}_n becomes unbounded and v_n is bounded, the output y_n must be also unbounded. This contradicts our original assumption and hence $H_u \equiv 0$. The original system h must be stable and the inputs $u^{(1)}, \dots, u^{(N)}$ can test stability in \mathcal{M}_{fd} .

(only-if part)

We now show that the two conditions for the inputs are also necessary to test the stability in \mathcal{M}_{fd} .

Suppose the first condition is not satisfied; consider an $\omega_0 \in [0, 2\pi]$ but $\omega_0 \notin \cup_{i=1}^N \Omega(u^{(i)})$. Consider the unstable system $h_n = \cos(n\omega_0)$. Lemma 5.9(b) implies that $u^{(i)} * e^{jn\omega_0}$ is bounded for all i . Since $u^{(i)} * h$ is the real part of $u^{(i)} * e^{jn\omega_0}$, it is also bounded for all i . Thus, the inputs cannot test stability in \mathcal{M}_{fd} . This shows that the first condition is necessary.

Now suppose that the second condition is not satisfied, so that there exists some $z_0 = r_0 e^{j\omega_0}$ ($0 < r_0 < 1$) which is a common zero in the open unit disk of the z -transforms of all the inputs; that is,

$$\sum_{k=0}^{\infty} u_k^{(i)} r_0^k e^{jk\omega_0} = 0, \quad \forall i \quad (5.44)$$

Since the inputs are real, their zeros occur as conjugate pairs, i.e.

$$\sum_{k=0}^{\infty} u_k^{(i)} r_0^k e^{-jk\omega_0} = 0 \quad \forall i \quad (5.45)$$

Now consider the unstable finite-dimensional system $h_n = r_0^{-n} \cos(n\omega_0)$. For each i, n ,

$$\begin{aligned} |(u^{(i)} * h)_n| &= \left| \sum_{k=0}^n u_k^{(i)} r_0^{-(n-k)} \cos(n-k)\omega_0 \right| \\ &= \left| \frac{1}{2} r_0^{-n} \sum_{k=0}^n u_k^{(i)} r_0^k (e^{j(n-k)\omega_0} + e^{-j(n-k)\omega_0}) \right| \\ &= \frac{1}{2} r_0^{-n} \left| e^{jn\omega_0} \left(\sum_{k=0}^n u_k^{(i)} r_0^k e^{-jk\omega_0} \right) + e^{-jn\omega_0} \left(\sum_{k=0}^n u_k^{(i)} r_0^k e^{jk\omega_0} \right) \right| \\ &= \frac{1}{2} r_0^{-n} \left| e^{jn\omega_0} \left(- \sum_{k=n+1}^{\infty} u_k^{(i)} r_0^k e^{-jk\omega_0} \right) + e^{-jn\omega_0} \left(- \sum_{k=n+1}^{\infty} u_k^{(i)} r_0^k e^{jk\omega_0} \right) \right| \\ &\leq r_0^{-n} \sum_{k=n+1}^{\infty} r_0^k = \frac{r_0}{1-r_0} \end{aligned}$$

Thus the output for each of the inputs is bounded. Hence the inputs $u^{(i)}$'s cannot test the stability in \mathcal{M}_{fd} . \square

We have the following corollary:

Corollary 5.10 \mathcal{M}_{fd} is testable for stability by a single input $u \in \overline{Bl}_\infty$ if and only if u excites at all frequencies and its z -transform has no zeros inside the open unit disk.

Neither the existence nor the non-existence of a bounded input having both the properties required by Corollary 5.10 has been established. However, bounded inputs which excite at all frequencies do exist. In fact, Lusin [12] has constructed a sequence which excites at all frequencies despite the fact that the sequence actually tends to 0. It will now be demonstrated that the input used in the proof of Theorem 4.4, containing all finite sequences of 1's and -1's, also excites at all frequencies.

Proposition 5.11 Let u be any sequence which contains all finite sequences of 1's and -1's. Then $\Omega(u) = [0, 2\pi]$.

Proof. Let ω_0 be an arbitrary frequency in $[0, 2\pi]$. Take any $M > 0$. The sum $\sum_k |\cos k\omega_0|$ is divergent, so we can find an integer L such that $\sum_{k=0}^L |\cos k\omega_0| > M$. By the definition of the sequence u , there exists an integer n_1 such that

$$(u_{n_1}, u_{n_1+1}, \dots, u_{n_1+L}) = (1, \operatorname{sgn}(\cos \omega_0), \operatorname{sgn}(\cos 2\omega_0), \dots, \operatorname{sgn}(\cos L\omega_0)) \quad (5.46)$$

Now,

$$\begin{aligned} \left| \sum_{k=n_1}^{n_1+L} u_k e^{-jk\omega_0} \right| &= \left| \sum_{k=n_1}^{n_1+L} \operatorname{sgn}(\cos(k - n_1)\omega_0) e^{-jk\omega_0} \right| \\ &= \left| \sum_{k=0}^L \operatorname{sgn}(\cos k\omega_0) e^{-jk\omega_0} \right| \\ &\geq \left| \sum_{k=0}^L \operatorname{sgn}(\cos k\omega_0) \cos k\omega_0 \right| > M \end{aligned}$$

This is true for every M , so $\limsup_{n \rightarrow \infty} \left| \sum_{k=0}^n u_k e^{-jk\omega_0} \right| = \infty \quad \square$

Proposition 5.11 leads to the following result:

Proposition 5.12 Two inputs can test stability in \mathcal{M}_{fd} .

Proof. Take $u^{(1)}$ as any input in \overline{Bl}_∞ containing all finite sequences of 1's and -1's, and $u^{(2)}$ as the unit impulse. Since $u^{(1)}$ excites at all frequencies and the z -transform of $u^{(2)}$ has no zeros, the result follows from Theorem 5.8. \square

This leads to the main result for the optimal asymptotic truncated l_1 error of \mathcal{M}_{fd} .

Theorem 5.13 *There are 2 inputs \mathbf{u} such that $E_{\infty}^t(\mathbf{u}, \mathcal{M}_{fd}, \delta) \leq 2\delta$*

Proof. Since \mathcal{M}_{fd} is testable for stability by two inputs, it follows from Theorem 4.4 that at most one additional input will be needed to achieve a worst-case diameter of 2δ . Recall from the proof of Theorem 4.4 that it is sufficient for the additional input to contain all finite sequences of 1's and -1's. But since this input can also be used as one of the inputs for testing stability, a saving of 1 input is obtained. Hence using the two inputs, we achieve $D(\mathbf{u}, \mathcal{M}_{fd}, \delta) = 2\delta$. Since \mathcal{M}_{fd} is σ -compact in the product topology (Proposition 5.5), the desired result now follows from Theorem 3.13. \square

The conclusion is therefore that \mathcal{M}_{fd} can be identified accurately in the limit. In particular, any model set containing finite-dimensional plants of bounded order can also be identified accurately in the limit.

C. All LTI Plants

A natural question to ask is whether there are any model sets larger than \mathcal{M}_{fd} that have small optimal asymptotic errors. In particular, what about \mathcal{X} , the set of all LTI plants? The results we have for this case are not complete.

Proposition 5.14 *No single input in $\overline{\mathcal{B}l}_{\infty}$ can test the stability of plants in \mathcal{X} . Hence, the optimal asymptotic error is infinite if only a single experiment is performed.*

Proof. Let $u \in \overline{\mathcal{B}l}_{\infty}$ be any bounded input. We shall find a plant $h \in \mathcal{X}$ which is unstable but yields an output satisfying $\|u * h\|_{\infty} \leq 1$.

Given u , the desired h can be constructed recursively as follows:

1) $h_0 = \frac{1}{u_0}$ (without loss of generality, u_0 can be assumed to be nonzero)

2) Given h_0, h_1, \dots, h_n , if

$$\left| -1 + \sum_{k=0}^n h_k u_{n+1-k} \right| < \left| 1 + \sum_{k=0}^n h_k u_{n+1-k} \right| \quad (5.47)$$

then define

$$h_{n+1} = -\frac{1}{u_0} \left(1 + \sum_{k=0}^n h_k u_{n+1-k} \right) \quad (5.48)$$

otherwise define

$$h_{n+1} = \frac{1}{u_0} \left(-1 + \sum_{k=0}^n h_k u_{n+1-k} \right) \quad (5.49)$$

Now, for every $n > 0$,

$$|h_{n+1}| = \max\left\{\left|\frac{1}{u_0}\left(1 + \sum_{k=0}^n h_k u_{n+1-k}\right)\right|, \left|\frac{1}{u_0}\left(-1 + \sum_{k=0}^n h_k u_{n+1-k}\right)\right|\right\} \geq \left|\frac{1}{u_0}\right|$$

so h is unstable. Also, $|u_0 h_0| = 1$ and

$$|(u * h)_{n+1}| = |h_{n+1} u_0 + \sum_{k=0}^n h_k u_{n+1-k}| = 1, \quad \forall n \geq 0$$

□

For the case of multiple experiments, we have the following necessary condition for the inputs to be able to test the stability of plants in \mathcal{X} . However, it is still an open problem where a finite number of bounded inputs can test the stability of plants in \mathcal{X} .

Proposition 5.15 *If the inputs $u^{(1)}, \dots, u^{(N)}$ can test the stability of plants in \mathcal{X} , then at least one of them must lie outside l_p for every $p < \infty$.*

Proof. Suppose that there exists some p , $1 \leq p < \infty$, such that for every i , $u^{(i)} \in l_p$. Let q satisfy $1/p + 1/q = 1$. Note that $1 < q \leq \infty$.

Consider the plant defined by $h_n = 1/(n+1)$. Since h is not in l_1 , it is unstable. However, it is in l_r for any $r > 1$. In particular, h is in l_q . For all i , $1 \leq i \leq N$, Holder's inequality yields $\|u^{(i)} * h\|_\infty \leq \|u^{(i)}\|_p \|h\|_q$. Hence the outputs of h to each of the inputs are bounded, although h is unstable. Therefore, the inputs cannot test the stability of plants in \mathcal{X} . □

6 Uniform Convergence of Estimates

In Section 3, it was shown that under some assumptions, the finite-horizon identification error will converge to the value of the infinite-horizon worst-case diameter, for all plants and all possible disturbances. However, nothing is said as to whether the convergence is *uniform*. Thus, although convergence is guaranteed for each plant, the rate of convergence may be arbitrarily slow for some plants in the model set. In this section, conditions on the model set will be presented that ensure uniform convergence of the estimates for all plants in the model set. This allows one to *a priori* determine the experiment length required to guarantee that *any* plant in the model set can be identified to a prescribed accuracy. In contrast to the σ -compactness condition that guarantees convergence, a compactness condition is used to guarantee uniform convergence.

Proposition 6.1 *If the model set \mathcal{M} is compact in the l_1 -topology, then there is an algorithm ϕ the estimates of which will converge uniformly to within $D(\mathbf{u}, \mathcal{M}, \delta)$ of the true plant; i.e. for all $\epsilon > 0$, there exists a time $T(\epsilon)$ such that for all $h \in \mathcal{M}$, $\|\mathbf{d}\|_\infty \leq \delta$,*

$$\|\phi(P_n \mathbf{u}, P_n(\mathbf{u} * h + \mathbf{d})) - h\|_1 \leq D(\mathbf{u}, \mathcal{M}, \delta) + \epsilon \quad \forall n > T(\epsilon)$$

Moreover, the algorithm does not require the knowledge of δ , the bound on the disturbances, to compute its estimates.

Proof. An algorithm ϕ is defined as follows: for each n ,

$$\phi(P_n \mathbf{u}, P_n \mathbf{y}) = \operatorname{argmin}_{g \in \mathcal{M}} \|P_n(\mathbf{u} * g - \mathbf{y})\|_\infty \quad (6.50)$$

The minimum must exist since \mathcal{M} is compact and $\|P_n(\mathbf{u} * g - \mathbf{y})\|_\infty$ is a continuous function of g in the l_1 -topology. Also note that computing this estimate does not require the knowledge of δ .

Now $\mathbf{y} = \mathbf{u} * h + \mathbf{d}$ for some true plant h and disturbance \mathbf{d} satisfying $\|\mathbf{d}\|_\infty \leq \delta$. By definition, the estimate at each time n satisfies

$$\|P_n(\mathbf{u} * \phi(P_n \mathbf{u}, P_n \mathbf{y}) - \mathbf{y})\|_\infty \leq \|P_n(\mathbf{u} * h - \mathbf{y})\|_\infty = \|P_n \mathbf{d}\|_\infty \leq \delta$$

and hence $\phi(P_n \mathbf{u}, P_n \mathbf{y}) \in S_n(\mathcal{M}, \mathbf{u}, \mathbf{y}, \delta)$ for each n , where S_n is the finite-horizon uncertainty set at time n . We shall use only this property of the estimates of ϕ to show that they uniformly converge.

Let $\epsilon > 0$. For each plant $h \in \mathcal{M}$, define

$$A(h, \epsilon) \equiv \{g \in \mathcal{M} : \|g - h\|_1 \geq D(\mathbf{u}, \mathcal{M}, \delta) + \epsilon\} \quad (6.51)$$

Also consider the number

$$T(\epsilon) \equiv \sup_{h \in \mathcal{M}} \sup_{g \in A(h, \epsilon)} T_{\mathbf{u}, \delta}(g - h) \quad (6.52)$$

where the function $T_{\mathbf{u}, \delta}$ has been defined in Section 3. $T(\epsilon)$ can be rewritten as $\sup_{h' \in B(\epsilon)} T_{\mathbf{u}, \delta}(h')$, where

$$B(\epsilon) = \{h' \in \mathcal{M} - \mathcal{M} : \|h'\|_1 \geq D(\mathbf{u}, \mathcal{M}, \delta) + \epsilon\}$$

and $\mathcal{M} - \mathcal{M} = \{g - h : g, h \in \mathcal{M}\}$. Since \mathcal{M} is compact and the subtraction operation is continuous with respect to the l_1 -topology, the set $\mathcal{M} - \mathcal{M}$ is also compact in the l_1 -topology. It follows that $B(\epsilon)$, a closed subset of $\mathcal{M} - \mathcal{M}$, is also compact. Now, $T_{\mathbf{u}, \delta}(h')$ is finite for all h' in $B(\epsilon)$, by Lemma 3.8. Hence, by Lemma 3.9, $T(\epsilon)$ is finite.

Now if $n > T(\epsilon)$, then for any plant $h \in \mathcal{M}$ and $\|\mathbf{d}\|_\infty \leq \delta$, the estimate $\hat{h}^{(n)}$ generated by the algorithm must lie in the uncertainty set $S_n(\mathcal{M}, \mathbf{u}, \mathbf{u} * h + \mathbf{d}, \delta)$. Hence, by Lemma 3.7, $T_{\mathbf{u}, \delta}(\hat{h}^{(n)} - h) > n > T(\epsilon)$. This implies

$$\|\hat{h}^{(n)} - h\|_1 < D(\mathbf{u}, \mathcal{M}, \delta) + \epsilon$$

Since this holds for all h and \mathbf{d} , the convergence is indeed uniform. \square

One can prove a similar result on the uniform convergence of estimates in the *truncated* l_1 norm.

Proposition 6.2 *If the model set \mathcal{M} is compact in the product topology, then there is an algorithm ϕ the estimates of which will converge uniformly, in the m -truncated l_1 seminorm, to within $D(\mathbf{u}, \mathcal{M}, \delta)$ of the true plant; i.e. for all $m \geq 0$ and $\epsilon > 0$, there exists $T(m, \epsilon)$ such that for all $h \in \mathcal{M}$, $\|\mathbf{d}\|_\infty \leq \delta$,*

$$\|P_m(\phi(P_n \mathbf{u}, P_n(\mathbf{u} * h + \mathbf{d})) - h)\|_1 \leq D(\mathbf{u}, \mathcal{M}, \delta) + \epsilon \quad \forall n > T(m, \epsilon)$$

Moreover, the algorithm does not require the knowledge of δ , the bound on the disturbances, to compute its estimates.

Equipped with these two results, we shall now analyse uniform convergence of estimates for the model sets \mathcal{M}_{stab} and \mathcal{M}_{fd} .

For \mathcal{M}_{stab} , it can at once be seen that although convergence to a small asymptotic error is possible, such convergence cannot be uniform.

Proposition 6.3 *Let ϕ be any algorithm and \mathbf{u} be any input. Then for every n and for every M , there exists an $h \in \mathcal{M}_{stab}$ such that*

$$\|\phi(P_n \mathbf{u}, P_n(\mathbf{u} * h)) - h\|_1 > M$$

Proof. This is clear because making n measurements gives no information on the part of the impulse response after time n , which can have arbitrarily large uncertainty in the l_1 norm. \square

To guarantee uniform convergence, we need to look at compact model sets.

Proposition 6.4 *Let $\mathcal{M} \subset \mathcal{M}_{stab}$ be a compact set (in the l_1 -topology) or a subset of a compact set in \mathcal{M}_{stab} . For the single input u^* which contains all finite sequences of 1's and -1's, there is an algorithm the estimates of which converge, uniformly for all $h \in \mathcal{M}$ and all $\|\mathbf{d}\|_\infty \leq \delta$, to an l_1 ball of radius 2δ around the true plant. Moreover the algorithm does not require the knowledge of the value of δ to compute its estimates.*

For example, any uniformly stable model set $\mathcal{M}_s(g)$ will admit uniform convergence in the l_1 norm.

Applying Proposition 6.2 to \mathcal{M}_{fd} , we have the following analogous result.

Proposition 6.5 *Let $\mathcal{M} \subset \mathcal{M}_{fd}$ be a compact set (in the product topology) or a subset of a compact set in \mathcal{M}_{fd} . Using two input, there is an algorithm the estimates of which converge, uniformly for all $h \in \mathcal{M}$ and all $\|\mathbf{d}\|_\infty \leq \delta$, to any truncated l_1 ball of radius 2δ around the true plant. Moreover the algorithm does not require the knowledge of the value of δ to compute its estimates.*

Thus, $\mathcal{M}_{fd}(p, q, K)$, the model set containing finite-dimensional plants with bounded parameters and order, will admit uniform convergence in the truncated l_1 norm.

7 Can Feedback Help?

The identification experiments in our problem formulation were assumed to be open-loop, i.e. the inputs are all chosen *a priori* and are not adapted based on the measured outputs. A natural question is whether *feedback* can improve the optimal worst-case asymptotic error. For model sets which are balanced, convex and σ -compact, feedback cannot in fact help much. This fact will be shown for the l_1 norm case; the truncated l_1 seminorm case follows similarly.

Suppose N closed-loop experiments are performed using a feedback policy μ , i.e. the inputs \mathbf{u} and the measured outputs \mathbf{y} are related by $\mathbf{u} = \mu(\mathbf{y})$. Here, μ is assumed to be strictly causal and satisfies $\mu(\mathbf{y}) \in \overline{B}_\infty^N$ for all measured outputs \mathbf{y} . Since \mathbf{u} and \mathbf{y} are vectors of input sequences, we take this to mean that the outputs can determine not only the future input sequence values within a *single* experiment but also which input sequences to use in *future* experiments.

Parallel to Definition 2.1, we have

Definition 7.1 *The worst-case asymptotic l_1 error using a feedback policy μ and an identification algorithm ϕ is*

$$e_\infty^1(\phi, \mathcal{M}, \mu, \delta) \equiv \sup_{h \in \mathcal{M}} \sup_{\|\mathbf{d}\|_\infty \leq \delta} \limsup_{n \rightarrow \infty} \|\phi(P_n(\mu(\mathbf{y})), P_n \mathbf{y}) - h\|_1$$

where given each h and \mathbf{d} , \mathbf{y} is the unique solution to

$$\mathbf{y} = \mu(\mathbf{y}) * h + \mathbf{d}$$

The optimal worst-case asymptotic l_1 error is

$$E_\infty^1(\mu, \mathcal{M}, \delta) \equiv \inf_{\phi} e_\infty^1(\phi, \mathcal{M}, \mu, \delta)$$

Parallel to Proposition 3.3, we have the following result

Proposition 7.2 *Let $D(\mu, \mathcal{M}, \delta)$ be the diameter of the largest possible infinite horizon uncertainty set given the feedback policy μ :*

$$D(\mu, \mathcal{M}, \delta) \equiv \sup_{h \in \mathcal{M}} \sup_{\|\mathbf{d}\|_\infty \leq \delta} \text{diam}(S_\infty(\mathcal{M}, \mu(\mathbf{y}), \mathbf{y}, \delta))$$

where \mathbf{y} is the unique solution to $\mathbf{y} = \mu(\mathbf{y}) * h + \mathbf{d}$. Then:

$$E_\infty^1(\mu, \mathcal{M}, \delta) \geq D(\mu, \mathcal{M}, \delta)/2$$

The following is the main result which establishes that feedback is not of much help in improving the optimal worst-case asymptotic error.

Proposition 7.3 *Let \mathcal{M} be a balanced, convex model set which is σ -compact in the l_1 topology. Suppose N closed-loop experiments are performed using feedback policy μ . Then there exist N inputs \mathbf{u} with which one can perform N open-loop experiments with an optimal worst-case asymptotic l_1 error of no more than twice that for the closed-loop experiments, i.e.*

$$E_\infty^1(\mathbf{u}, \mathcal{M}, \delta) \leq 2E_\infty^1(\mu, \mathcal{M}, \delta) \quad \forall \delta \geq 0$$

Proof. Let $\mathbf{u}^0 \equiv \mu(0)$ be the input given by the feedback policy when there is zero output. (This can occur when the true plant and the disturbances are all zero.) We have, for every $\delta \geq 0$,

$$\begin{aligned} 2E_\infty^1(\mu, \mathcal{M}, \delta) &\geq D(\mu, \mathcal{M}, \delta) \\ &\geq \text{diam}(S_\infty(\mathcal{M}, \mathbf{u}^0, 0, \delta)) \\ &= D(\mathbf{u}^0, \mathcal{M}, \delta) \quad \text{by Prop. 4.2} \\ &\geq E_\infty^1(\mathbf{u}^0, \mathcal{M}, \delta) \end{aligned}$$

where the last inequality follows from Theorem 3.13. \square

This result is related to a general result in information-based complexity theory that “adaptive linear information” is no more powerful than “non-adaptive linear information” in linear problems, when the problem space is balanced and convex [22].

8 Conclusions

In this paper, we have approached the problem of analyzing the intrinsic limitations of identification by considering the optimal worst-case asymptotic error achievable using any input and any identification algorithm. This gives an intrinsic measure of the difficulty of identification, given the *a priori* knowledge (model set and disturbance class) and the constraints on the allowable experiments (input class).

The analysis is performed in two steps. First, for fixed inputs, a lower bound on the error of any identification algorithm is expressed in terms of the diameter of the worst-case infinite-horizon uncertainty set, and it was shown that under some compactness conditions on the model set, there exist algorithms which achieve this bound asymptotically. Second, characterization of inputs which minimize this worst-case diameter is given. Applying these general results, it was shown that large model sets, such as the class of all stable systems and the class of all finite-dimensional systems, can be identified very accurately in the limit.

To obtain concrete results, we have restricted to using the l_1 norm as the error (or cost) measure, and assumed that both the inputs and the disturbances are constrained to lie in l_∞ balls. However, the framework can also be applied to different choices of error measures and disturbance and input classes. For example, one can consider using the H_∞ metric as an error measure, and impose energy instead of magnitude bounds on the disturbances and inputs. Another possible generalization is to use a weighted norm on the disturbance to reflect the prior knowledge that the energy is concentrated in some frequency bands. It can be shown that, under very mild assumptions on the norms, the consistency results proven in Sections 3 and 6 in fact holds for arbitrary error and disturbance norms. In general, if the model set is σ -compact in the topology generated by the error norm, then the infinite-horizon error can be achieved by finite-horizon experiments in the limit; furthermore, the convergence is guaranteed to be uniform if the model set is compact.

There are three major directions to be pursued in future research:

Further Work in Identification

This includes

- a) Investigate possible generalizations to MIMO and continuous-time systems.
- b) Analyze the problem of having input in addition to output noise.
- c) Study model sets that incorporate some qualitative information. For example, if we are told that the impulse response is unimodal, how should we choose the optimal input for the identification

experiment? Note that such a model set is not convex.

d) Study the extent to which randomly chosen inputs can improve the potential for accurate identification. Such improvement would correspond to the fact that mixed strategies can outperform pure strategies in game theory.

e) Study the optimal rate of convergence of the finite-horizon estimates. That is, how long does one have to perform the experiment to guarantee a desired accuracy for the estimate?

f) Explore the relationship of this work with other measures of intrinsic difficulty of identification problems, such as using metric entropy of model sets [26].

Relationship to Robust Control

The results presented in this paper have implications on robust control. From the point of view of control, a model is accurate if it can be used to design a controller which will perform well with the actual plant. Accurate identification in the l_1 norm satisfies this criterion, for the following reason: Given any stabilizing controller, the closed-loop system, viewed as a map on bounded signals, varies continuously as a function of the open-loop plant, in the l_1 topology. Hence if the identified model is close to the actual plant in the l_1 norm, the desired closed-loop performance will not differ too much from the actual closed-loop performance. Hence, our result indicates that the set of all stable plants can be identified accurately, in the limit, for the purpose of control.

Accurate identification in the truncated l_1 seminorm, on the other hand, does not necessarily have this property. One has to compare the product topology to the *graph* topology, which is the weakest topology in which the closed-loop system varies continuously as a function of the open-loop system [25]. We expect that although the two topologies are in general different, they could in fact be *equivalent* when restricted to certain model sets. For such model sets, accurate identification in the truncated l_1 seminorm is sufficient for the purpose of control.

At a broader level, this work sheds a new light on the robust control problem. Given the prior information, the outcome of any identification scheme is an uncertainty set in which the actual plant lies. The challenge is to come up with a control strategy that stabilizes every plant in the set and at the same time meet specific performance objectives. This requires a complete characterization of this uncertainty set in a simple fashion and then the development of new robust control techniques to handle such a set. This will be a future area of research.

Relationship to Adaptive Control

Finally, we believe that the framework and the machinery developed in this paper can be

extended to study the fundamental limitations and capabilities of adaptive control. Here, given a model set \mathcal{M} (possibly containing unstable plants), we pose the following questions:

- a) Is there an adaptive controller which stabilizes all plants in \mathcal{M} ?
- b) Is there an adaptive controller whose asymptotic performance is comparable to the closed-loop performance that could be attained if the plant were known exactly?

Such questions have been insufficiently studied at this level of generality. We believe a further development of our methodology will significantly enhance our understanding of these problems.

9 Acknowledgements

We would like to thank Sanjeev Kulkarni and Emre Teletar for the very helpful discussions.

References

- [1] A. Blumer, A. Ehrenfeucht, D. Haussler, M. Warmuth, "Occam's Razor", *Information Processing Letters* 24, pp.377-380, 1987.
- [2] M.A. Dahleh and J. B. Pearson, Jr., "Optimal rejection of persistent disturbances, robust stability, and mixed sensitivity minimization", *IEEE Trans. Automatic Control*, AC-33, No.8, pp.722-731, 1988.
- [3] M.A. Dahleh and Y. Ohta, "A necessary and sufficient condition for BIBO stability", *Systems and Control Letters*, 11, pp. 271-275, 1988.
- [4] E. Fogel and Y. F. Huang, "On the value of information in system identification—bounded noise case", *Automatica*, vol.18, no.2, pp.229-238, 1982.
- [5] G.C. Goodwin, "Experiment design in system identification" in *Encyclopedia of Systems and Control* (M. Singh, Ed.), Pergamon Press, 1987.
- [6] A.J. Helmicki, C.A. Jacobson and C.N. Nett, "Identification in H_∞ : A robust convergent nonlinear algorithm", Proceedings of the 1989 International Symposium on the Mathematical Theory of Networks and System, 1989.
- [7] A.J. Helmicki, C.A. Jacobson and C.N. Nett, "Identification in H_∞ : Linear Algorithms", Proceedings of the 1990 American Control Conference, pp 2418-2423.
- [8] M. Khammash and J.B. Pearson, "Stability and Performance Robustness in l_1 optimal control systems", accepted for publication in *IEEE Trans. Automatic Control*
- [9] J.M. Krause, G. Stein, P.P. Khargonekar, "Robust Performance of Adaptive Controllers with General Uncertainty Structure", Proceedings of the 29th Conference on Decision and Control, pp. 3168-3175, 1990.
- [10] R. Lozano-Leal and R. Ortega, "Reformulation of the parameter identification problem for systems with bounded disturbances", *Automatica*, vol.23, no.2, pp.247-251, 1987.
- [11] M.K. Lau, R.L. Kosut, S. Boyd, "Parameter Set Estimation of Systems with Uncertain Nonparametric Dynamics and Disturbances", Proceedings of the 29th Conference on Decision and Control, pp. 3162-3167, 1990.

- [12] A. Markushevich, *Theory of Functions of A Complex Variable*, Volume I, Prentice Hall, 1965.
- [13] R.K. Mehra, "Optimal input signals for parameter estimation in dynamic systems—A survey and new results", *IEEE Trans. Automatic Control*, vol AC-19, pp.753-768, 1974.
- [14] C.A. Michelli and T.J. Rivlin, "A survey of optimal recovery" in *Optimal Estimation in Approximation Theory* (C.A. Michelli and T.J. Rivlin, Eds), Plenum, New York, 1977.
- [15] M. Milanese and G. Belforte, "Estimation theory and uncertainty intervals evaluation in the presence of unknown but bounded errors: Linear families of models and estimators", *IEEE Trans. Automatic Control*, AC-27, pp.408-414, 1982.
- [16] M. Milanese and R. Tempo, "Optimal algorithm theory for robust estimation and prediction", *IEEE Trans. Automatic Control*, AC-30, pp. 730-738, 1985.
- [17] M. Milanese, R. Tempo and A. Vicino, "Strongly optimal algorithms and optimal information in estimation problems", *Journal of Complexity*, vol.2, pp.78-94, 1986.
- [18] J. Munkres, *Topology - A First Course*, Prentice Hall, 1975.
- [19] J.P. Norton, "Identification and application of bounded-parameter models", *Automatica*, vol.23, no.4, pp.497-507, 1987.
- [20] J. Pearl, "On the Connection between the Complexity and Credibility of Inferred Models", *Internat. J. General Systems*, 4, , 255-264, 1978.
- [21] L. Pronzato and E. Walter, "Experiment design in bounded-error context: Comparison with D-optimality", *Automatica*, vol.25, no.3, pp.383-391, 1989.
- [22] J.F. Traub and H. Wozniakowski, *A General Theory of Optimal Algorithms*, Academic Press, New York, 1980.
- [23] J.F. Traub, G. Wasilkowski and H. Wazniakowski, *Information, Uncertainty, Complexity*, Addison-Wesley, Reading, Mass., 1983.
- [24] L. Valiant, "A Theory of the Learnable", *Comm. of the ACM*, 27(11), 1984.
- [25] M. Vidyasagar, *Control System Synthesis*, MIT Press, Cambridge, Mass., 1985.
- [26] G. Zames, "On the metric complexity of casual linear systems: ϵ -entropy and ϵ -dimension for continuous-time", *IEEE Trans. on Automatic Control*, Vol. 24, April 1979.
- [27] M. Zarrop, *Optimal Experimental Design for Dynamic System Identification*, Springer-Verlag, New York, 1979.