Optimal Asymptotic Identification Under Bounded Disturbances

David N. C. Tse, *Student Member, IEEE*, Munther A. Dahleh, *Member, IEEE*, and John N. Tsitsiklis, *Member, IEEE*

***Abstract*—This paper investigates the intrinsic limitation of worst-case 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. First, it is shown that under some topological conditions on the model set, there is an identification algorithm which is asymptotically optimal for any input. Characterization of the optimal asymptotic error as a function of the inputs is also obtained. These results hold for any error metric and disturbance norm. Second, these general results are applied to three specific identification problems: identification of stable systems in the l_1 norm, identification of stable rational systems in the H_∞ norm, and identification of unstable rational systems in the gap metric. For each of these problems, the general characterization of optimal asymptotic error is used to find near-optimal inputs to minimize the error.**

I. Introduction

RECENTLY, there has been a growing line of work with the common theme that system identification should be performed so that the *worst-case* error of the resulting model is small in a metric compatible with robust control [8]–[10], [26], [37]. This paper addresses the questions of asymptotically optimal identification algorithms and experiment designs from this point of view. Our emphasis is less on finding efficient algorithms and more on finding the *fundamental limitations* in identification accuracy achievable by *any* identification algorithm in the limit of observing more and more data corrupted by nonstochastic noise. Thus, this work is in the flavor of the questions posed by Zames [41].

We will deal exclusively with discrete-time, single–input single–output linear time-invariant systems. In this formulation, the unknown plant is *a priori* known to be in a certain subset $\mathfrak{M}$ of the space of all LTI systems; this subset will be called a *model set* $\mathfrak{M}$. The model set is endowed with a general metric ρ which can be any uncertainty measure suitable for designing robust controllers. To identify the plant, one is allowed to perform one or more finite but arbitrarily long experiments using input sequences chosen from a given input set $\mathfrak{U}$. (Typically, $\mathfrak{U}$ is some norm-bounded set.) The measured outputs are corrupted with additive disturbance sequences which are bounded in an l_p norm $\|\cdot\|_p$ but can otherwise be arbitrary. The problem is to analyze the *smallest* worst-case error, over all plants in $\mathfrak{M}$ and all admissible disturbances, achievable by using any inputs from $\mathfrak{U}$ and any identification algorithm to estimate the plant from arbitrarily long but finite data records (i.e., asymptotic error). 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.

A natural framework to study worst-case identification is provided by information-based complexity theory [21], [35], [36]. 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. Our work, like many others in worst-case identification, has employed some of the basic concepts of this theory, but the key results we derived are completely new.

Although mainstream system identification research adopts stochastic models for the noise, there is a line of work which deals with worst-case identification under bounded disturbances [5], [16], [22]–[24], [28], [32], [15]. More recently, specific identification algorithms are proposed in [8]–[10], [26] for worst-case identification in the H_∞ metric from noisy frequency response data and in [12], [25] for identification in the l_1 metric from time series data. In contrast to these works, we deal with general aspects of optimal worst-case asymptotic identification in a general error metric. Moreover, the issue of optimal experiment design, although considered in stochastic system identification (e.g., [7], [20], [43]), has not been satisfactorily addressed in the worst-case setting. Issues of complexity and tradeoffs between the length of experiments and accuracy has been recently reported in [3], [13], [18], [31].

Manuscript received October 23, 1991; October 25, 1992. Paper recommended by Past Associate Editor, T. T. Georgiou. This work was supported in part by an NSERC fellowship from the government of Canada, by the NSF under Grants ECS-8552419, ECS-9157306, and ECS-8552419, by the AFOSR under Grants AFOSR-91-0368 and AFOSR-91-0368, and by the ARO under Grants DAAL03-86-K-0171 and DAAL03-86-K-0171.

The authors are with the Laboratory for Information and Decision Systems, M.I.T., Cambridge, MA 02139.

IEEE Log Number 9210358.

0018-9286/93$03.00

The contributions of this paper are two-folded. 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. At a more specific level, concrete results are obtained by applying the general framework to three specific identification problems: identification of stable systems in the l_1 and H_∞ metrics, and identification of unstable systems in the gap metric. In all these problems, the required topological conditions for consistency are verified and the infinite-horizon problem is analyzed to find good input designs.

The organization of the paper is as follows. In Section II, the identification problem is formulated and the optimal worst-case asymptotic error achievable by any identification algorithm is defined. In Section III, we present consistency results establishing infinite-horizon experiments as limits of finite-horizon ones. In Section IV, the general results developed are applied to analyze three specific identification problems. Section V contains our conclusion.

II. Problem Formulation

Let $\mathfrak{X}$ be the class of all causal, single–input single–output, linear time-invariant, discrete-time systems. We identify $\mathfrak{X}$ with the space of all one-sided real-valued sequences, $\boldsymbol{R}^{\omega}$. Let $\mathfrak{M} \subset \mathfrak{X}$ be the *model set* which is assumed to contain the unknown plant h to be identified. The set $\mathfrak{M}$ captures the experimenter's *a priori* knowledge about h. Some examples of $\mathfrak{M}$ are the set of all stable systems, the set of stable systems with a bound on the decay rate, the set of all finite-dimensional systems with a bound on the order, etc. Also given is an input set $\mathfrak{U}$ which contains all the input sequences that can be used in the identification experiments. Typically, $\mathfrak{U}$ is a norm-bounded set, to reflect physical limitations, power restrictions, safety, or to maintain the validity of the linear model of the plant.

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

$$y = h * u + d \tag{2.1}$$

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; $h = (h_0, h_1, h_2, \cdots,$ etc.). The disturbance d is assumed to be bounded in a given norm, $\|d\|_p \le \delta$ for some known δ, but can otherwise be arbitrary. The disturbance may arise from actual measurement noise, such as quantization, or it may reflect nonlinearities and time-variation of the plant. In the latter case, the true plant is actually nonlinear and time varying but is assumed to be approximated well at the operating range by an LTI component, which is the object of identification.

One point to note is that we assume that the system is initially at rest before an experiment is started. Having an unknown nonzero initial condition is equivalent to having an additional, unknown, additive disturbance $u^- * h$, where u^- is the (unknown) input before time $t = 0$. If the model set $\mathfrak{M}$ is bounded in the operator norm from the input space to the disturbance space, then $u^- * h$ is bounded if u^- is, and this additional uncertainty can be accounted for by grouping into the original additive disturbance term. If this is not the case, however, then the problem cannot be treated in the present framework.

Now suppose N such independent experiments are performed. The question whether more than one input is needed to identify plants in a given model set will be addressed. We then have:

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

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

$$\boldsymbol{y} = \boldsymbol{u} * h + \boldsymbol{d} \qquad \|\boldsymbol{d}\|_p \equiv \max_i \|d^{(i)}\|_p \le \delta \tag{2.3}$$

where $\boldsymbol{y} = [y^{(1)}, \cdots, y^{(N)}]$, $\boldsymbol{u} = [u^{(1)}, \cdots, u^{(N)}]$, and $\boldsymbol{d} = [d^{(1)}, \cdots, d^{(N)}]$ 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 $\boldsymbol{u}$ is in $\mathfrak{U}^N$.

An *identification algorithm* is a mapping ϕ which generates, at each time instant n, an estimate $\hat{h}^{(n)} \equiv \phi(P_n \boldsymbol{u}, P_n \boldsymbol{y}) \in \mathfrak{X}$ of the unknown plant h, given the input and output sequences in the experiments. Here, P_n is the truncation operator, defined by $P_n x = (x_0, x_1, \cdots, x_n)$ 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. Generally, we will assume that the algorithm has access to what the model set $\mathfrak{M}$ is and also the value of δ, the bound on the disturbance. In the terminology of Helmicki *et al.* [12], the algorithm is *tuned*. However, in some cases, we will be able to give stronger results using algorithms which are untuned to the value of δ.

Also given is an extended metric $\rho(\cdot,\cdot)$ on $\mathfrak{X}$, $\rho: \mathfrak{X} \times \mathfrak{X} \to \boldsymbol{R} \cup \{\infty\}$, which evaluates the accuracy of $\hat{h}^{(n)}$ as an estimate of h.

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

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

$$\limsup_{n\to\infty} \rho(\phi(P_n u, P_n(u * h + d)), h) \leq r.$$

Equivalently,

$$e_\infty(\phi, \mathfrak{M}, u, \delta) \equiv \sup_{h\in\mathfrak{M}} \sup_{\|d\|_\infty \leq \delta} \limsup_{n\to\infty} \rho(\phi(P_n u, P_n(u * h + d)), h).$$

According to this definition, no matter what the true plant and the disturbances are, the plant can be eventually approximated to within $e_\infty(\phi, \mathfrak{M}, u, \delta)$, 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.

In the above definition of the worst-case asymptotic error, although convergence of the estimates to within $e_\infty(\phi, \mathfrak{M}, u, \delta)$ is guaranteed for all admissible plants and disturbance sequences, the *rate* of convergence may be arbitrarily slow for some plants and some disturbances. The worst-case asymptotic error is said to be *uniform* if the rate of convergence is uniform over all admissible plants and disturbance sequences. If the convergence is uniform, the worst-case asymptotic error defined above is the same as the limit of the worst-case error taken at each finite time n, i.e.,

$$\sup_{h\in\mathfrak{M}} \sup_{\|d\|_\infty \leq \delta} \limsup_{n\to\infty} \rho(\phi(P_n u, P_n(u * h + d)), h) = \limsup_{n\to\infty} \sup_{h\in\mathfrak{M}} \sup_{\|d\|_\infty \leq \delta} \rho(\phi(P_n u, P_n(u * h + d)), h)$$

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. It is the notion of convergence considered by Helmicki *et al.* in their framework [11].

Demanding uniform convergence is too restrictive a formulation for a general theory of fundamental limitations of worst-case identification. Although such uniform convergence is certainly desirable, it is impossible to achieve for many interesting model sets. In fact, for many inherently infinite-dimensional model sets, the worst-case error at each finite time is always infinite, while the worst-case asymptotic error can be made small using an appropriate identification algorithm and inputs. Our formulation thus allows us to discuss optimal worst-case identification and optimal inputs for a much broader class of model sets. Besides, in some applications of identification, such as adaptive control, uniform convergence of estimates is not necessary to fulfill the desired objectives. However, because of the special importance of uniform convergence, we will give additional conditions on the model set for this to take place. It will be seen that these conditions are quite strong and essentially require the model set to be finite-dimensional. It is worthwhile to note that the model set considered in [8], [9] satisfies these conditions.

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

$$E_\infty(u, \mathfrak{M}, \delta) \equiv \inf_{\phi} e_\infty(\phi, \mathfrak{M}, u, \delta).$$

Any algorithm for which the infimum is attained is said to be *asymptotically optimal.* We will obtain a general characterization of the asymptotically optimal algorithms and the resulting optimal worst-case asymptotic error, for given inputs u. For specific problems, we will find conditions on the inputs u to make this optimal worst-case asymptotic error small.

It should be noted that the asymptotically optimal algorithms to be derived are valid for arbitrary inputs u. This allows the complete separation of the problem of devising optimal algorithms and the problem of designing optimal inputs. This is particularly important when there is no complete control over the choice of the inputs into the plants, such as in closed-loop experiments or in adaptive control. In these problem, this "separation principle" facilitates the derivation of necessary conditions on the input signals for accurate identification to take place.

We would also like to point out that there are some recent asymptotic optimality results in the general information-based complexity framework [14]. However, their notion of optimality is that of the *rate of convergence* of the worst-case error for any *fixed* problem element, and their results only make sense if the error converges to zero. In contrast, in the worst-case identification problem we are dealing with, the error does not typically converge to zero, and our notion of optimality is that of the nonzero limit supremum of the error.

III. Asymptotically Optimal Identification

In this section, the inputs will be assumed to be fixed. The characterization of asymptotically optimal algorithms and optimal worst-case asymptotic error is in terms of the important notion of the *uncertainty set*, an important notion in information-based complexity theory.

Definition 3.1: Let u and 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

$$S_n(\mathfrak{M}, u, y, \delta) = \{g \in \mathfrak{M}: \|P_n(u * g - y)\|_p \leq \delta\}$$

and the infinite-horizon uncertainty set is

$$S_\infty(\mathfrak{M}, u, y, \delta) = \{g \in \mathfrak{M}: \|u * g - y\|_p \leq \delta\}.$$

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.

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

$$\operatorname{diam}(A) = \sup_{g,h \in A} \rho(g,h),$$

$$\operatorname{rad}(A) = \inf_{g \in \mathfrak{X}} \sup_{h \in A} \rho(g,h).$$

Note that $\operatorname{diam}(A)/2 \leq \operatorname{rad}(A) \leq \operatorname{diam}(A)$. We shall now define two important quantities.

Definition 3.2: Given a choice of the inputs $\boldsymbol{u}$, define the infinite-horizon diameter of information $D(\boldsymbol{u}, \mathfrak{M}, \delta)$ and radius of information $R(\boldsymbol{u}, \mathfrak{M}, \delta)$ to be respectively the diameter and radius of the largest possible uncertainty set:

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

$$R(\boldsymbol{u}, \mathfrak{M}, \delta) \equiv \sup_{h \in \mathfrak{M}} \sup_{\|\boldsymbol{d}\|_\infty \leq \delta} \operatorname{rad}(S_\infty(\mathfrak{M}, \boldsymbol{u}, \boldsymbol{u} * h + \boldsymbol{d}, \delta)).$$

In information-based complexity terminology, these quantities correspond to the *diameter* and *radius of information* for the infinite-horizon problem where the information available is the entire infinite output sequence. The quantity $D(\boldsymbol{u}, \mathfrak{M}, \delta)$ is the largest distance between two plants for which there are admissible disturbances such that the plants give exactly the same outputs. It turns out that it is precisely this quantity that characterizes the optimal worst-case asymptotic errors. First we show that half the infinite-horizon diameter of information is a lower bound to the optimal asymptotic error.

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

$$e_\infty(\phi, \mathfrak{M}, \boldsymbol{u}, \delta) \geq D(\boldsymbol{u}, \mathfrak{M}, \delta)/2$$

for any algorithm ϕ.

Proof: Let ψ be an algorithm for the infinite-horizon problem, i.e., given the entire input and output sequences, ψ generates an estimate for the plant. The worst-case error achieved by this algorithm is:

$$\sup_{h \in \mathfrak{M}} \sup_{\|\boldsymbol{d}\|_\infty \leq \delta} \rho(\psi(\boldsymbol{u}, \boldsymbol{u} * h + \boldsymbol{d}), h)$$

and the infinite-horizon optimal worst-case error achievable by any algorithm is

$$\inf_{\psi} \sup_{h \in \mathfrak{M}} \sup_{\|\boldsymbol{d}\|_\infty \leq \delta} \rho(\psi(\boldsymbol{u}, \boldsymbol{u} * h + \boldsymbol{d}), h). \tag{3.4}$$

One should note that while the algorithms allowed in this infinite-horizon problem have access to the entire infinite input–output sequences, the algorithms for the asymptotic problem have access to only finite but arbitrarily long portions. Consequently, the infinite-horizon optimal worst-case error lower bounds the optimal asymptotic error $E_\infty(\boldsymbol{u}, \mathfrak{M}, \delta)$. On the other hand, by a central result in information-based complexity theory [35], this infinite-horizon optimal error is given by the infinite-horizon radius of information $R(\boldsymbol{u}, \mathfrak{M}, \delta)$, which in turn is lower bounded by half the diameter of information $D(\boldsymbol{u}, \mathfrak{M}, \delta)$. Hence, the result follows. □

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 $D(\boldsymbol{u}, \mathfrak{M}, \delta)$ 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 $D(\boldsymbol{u}, \mathfrak{M}, \delta)$ 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 that $D(\boldsymbol{u}, \mathfrak{M}, \delta)$), 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 $T(h') < n$ 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 a 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 [17], 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 an unknown-but-bounded noise setting. In fact, it will now be shown that a compactness condition on the model set will guarantee consistency.

The following theorem shows that, under a σ-compactness assumption on $\mathfrak{M}$, $D(\boldsymbol{u}, \mathfrak{M}, \delta)$ is an upper bound for the optimal asymptotic 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 error is reduced to the study of $D(\boldsymbol{u}, \mathfrak{M}, \delta)$, if we ignore this factor of 2.

Theorem 3.4: Suppose that the model set $\mathfrak{M}$ is σ-compact in the ρ-topology, $\mathfrak{M} = \bigcup_i \mathfrak{M}_i$, $M_i \subset M_{i+1}$ $\forall_i$, $\mathfrak{M}_i$ compact and on each $\mathfrak{M}_i$, convergence in the ρ-topology implies component-wise convergence of the impulse response. Then there is an identification algorithm ϕ^* such that $e_\infty(\phi^*, \mathfrak{M}, \boldsymbol{u}, \delta) \leq D(\boldsymbol{u}, \mathfrak{M}, \delta)$ for all $\boldsymbol{u}$ and $\delta \geq 0$.

It should be noted that by an elementary result in information-based complexity theory, the optimal worst-case error achievable when the algorithm has *full* access to the entire *infinite* input–output sequences is also bounded between the infinite-horizon diameter of information and half the diameter of information. Our two results (Proposition 3.3 and Theorem 3.4) are of an entirely different nature: they assert that the optimal worst-case asymptotic error achievable when the algorithm has access to *finite* but arbitrarily long data records also satisfies the same bounds. The assumed topological conditions are crucial for the validity of Theorem 3.4.

Before proving Theorem 3.4, we need one more definition and a few lemmas.

Definition 3.5: For given inputs $\boldsymbol{u}$ and bound δ on disturbances, and $g, h \in \mathfrak{X}$, define $T_{u,\delta}(g,h)$ to be the smallest integer k such that $\|P_k(\boldsymbol{u} *(g-h))\|_p > 2\delta$. If no such k exists, then $T_{u,\delta}(g,h)$ is infinite.

Lemma 3.6: For any two plants $g, h \in \mathfrak{M}$, $T_{u,\delta}(g,h)$ is the smallest k such that there is no output y with g and h in the same uncertainty set $S_k(\mathfrak{M}, \boldsymbol{u}, \boldsymbol{y}, \delta)$.

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

Thus, given two plants g and h, $T_{u,\delta}(g,h)$ 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 3.7: Let $g, h \in \mathfrak{M}$. If $\rho(g,h) > D(\boldsymbol{u}, \mathfrak{M}, \delta)$, then $T_{u,\delta}(g,h) < \infty$.

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

The desired topological condition involves the topology of component-wise convergence of sequences, or the so-called product topology [27].

Lemma 3.8: Fix the inputs $\boldsymbol{u} \in \overline{Bl}_\infty^N$ and $\delta > 0$. Let $A \subset \mathfrak{M} \times \mathfrak{M}$ be compact in the product topology, and suppose $T_{u,\delta}(g,h)$ is finite for every $(g,h) \in A$. Then $\sup_{(g,h) \in A} T_{u,\delta}(g,h)$ is also finite.

Proof: Suppose $\sup_{(g,h) \in A} T_{u,\delta}(g,h) = \infty$. Then there exists a sequence of plants $(g^{(i)}, h^{(i)})$ in A such that $\lim_{i\to\infty} T_{u,\delta}(g^{(i)}, h^{(i)}) = \infty$; furthermore, the sequence can be assumed to converge (in the product topology) to a pair of plant $(g^*, h^*) \in A$ since A is compact. Let $n^* \equiv T_{u,\delta}(g^*, h^*) < \infty$. By definition, $\|P_{n^*}(\boldsymbol{u} *(g^* - h^*))\|_p > 2\delta$. Since the norm of a sequence is a continuous function of finitely many of its components, it follows that $\|P_{n^*}(\boldsymbol{u} *(g-h))\|_p$ is a continuous function of (g,h) in the product topology. Hence, there exists a ball B (in the product topology) around (g^*, h^*) such that for every $(g', h') \in B$, $\|P_{n^*}(\boldsymbol{u} *(g'-h'))\|_p > 2\delta$, i.e., $T_{u,\delta}(g', h') \leq n^*$ for every $(g', h') \in B$. But this contradicts the fact that $\lim_{i\to\infty} T_{u,\delta}(g^{(i)}, h^{(i)}) = \infty$ since $(g^{(i)}, h^{(i)}) \to (g^*, h^*)$. Hence, it can be concluded that $\sup_{(g,h) \in A} T_{u,\delta}(g,h)$ is in fact finite. □

Basically, this lemma says that if each plant in the compact set A can be eventually ruled out as the true plant, there is a finite time after which all of them can be simultaneously ruled out.

Now we are in a position to prove Theorem 3.4.

Proof: 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 $S_n \cap \mathfrak{M}_k$, where S_n is the uncertainty set after observing the output data until time n, and k is the least integer i such that $S_n \cap \mathfrak{M}_i$ is nonempty. We claim that this algorithm will have an asymptotic error of at most $D(\boldsymbol{u}, M, \delta)$ for all inputs $\boldsymbol{u}$ and $\delta > 0$.

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

$$A(h,\epsilon) \equiv \{g \in \mathfrak{M}_h : \rho(g,h) \geq D(\boldsymbol{u}, \mathfrak{M}, \delta) + \epsilon\} \tag{3.5}$$

and the number

$$T(h,\epsilon) \equiv \sup_{g \in A(h,\epsilon)} T_{u,\delta}(g,h). \tag{3.6}$$

Since $A(h,\epsilon)$ is a closed subset of $\mathfrak{M}_h$ (with respect to the ρ-topology), it is also compact in the ρ-topology. Since the ρ-topology is finer than the product topology in $\mathfrak{M}_h$, $A(h,\epsilon)$ is also compact in the product topology. By Lemma 3.7, $T_{u,\delta}(g,h)$ is finite for all $(g,h) \in A(h,\epsilon)$. Hence, by Lemma 3.8, $T(h,\epsilon)$ 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 $S_n \cap \mathfrak{M}_k$ is nonempty, $\hat{h}^{(n)}$ is guaranteed to be in M_h for all n. (This is because $S_n \cap \mathfrak{M}_h$ is nonempty: it contains the true plant h.) Also $\hat{h}^{(n)}$ is in the uncertainty set S_n and by Lemma 3.6, $T_{u,\delta}(\hat{h}^{(n)}, h) > n$. If we now take any $n > T(h,\epsilon)$, we have $T_{u,\delta}(\hat{h}^{(n)}, h) > T(h,\epsilon)$ so $\hat{h}^{(n)}$ is not in $A(h,\epsilon)$. But $\hat{h}^{(n)}$ is in $\mathfrak{M}_h$, so it follows that $\rho(\hat{h}^{(n)}, h) < D(\boldsymbol{u}, \mathfrak{M}, \delta) + \epsilon$.

Since ϵ is arbitrary, it can now be concluded that

$$\limsup_{n\to\infty} \rho(h^{(n)}, h) \leq D(\boldsymbol{u}, \mathfrak{M}, \delta)$$

completing the proof. □

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, $\mathfrak{M} = \bigcup_i \mathfrak{M}_i$. Given this nested structure, plants in the smaller $\mathfrak{M}_i$'s are considered to be simpler than those in larger $\mathfrak{M}_i$. 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], [30].

In contrast to the σ-compactness condition that guarantees convergence, a stronger compactness condition guarantees uniform convergence.

Proposition 3.9: Suppose convergence in the ρ-topology on $\mathfrak{M}$ implies component-wise convergence of the impulse response. If the model set $\mathfrak{M}$ is compact in the ρ-topology, then there is an algorithm ϕ the estimates of which will converge uniformly to within $D(\boldsymbol{u}, \mathfrak{M}, \delta)$ of the true plant; i.e., for all $\epsilon > 0$, there exists a time $T(\epsilon)$ such that for all $h \in \mathfrak{M}$, $\|\boldsymbol{d}\|_p \le \delta$,

$$\rho(\phi(P_n \boldsymbol{u}, P_n(\boldsymbol{u} * h + \boldsymbol{d})), h) \le D(\boldsymbol{u}, \mathfrak{M}, \delta) + \epsilon \qquad \forall n > T(\epsilon).$$

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,

$$\phi(P_n \boldsymbol{u}, P_n \boldsymbol{y}) = \operatorname*{argmin}_{g \in \mathfrak{M}} \|P_n(\boldsymbol{u} * g - \boldsymbol{y})\|_p. \tag{3.7}$$

The minimum must exist since $\mathfrak{M}$ is compact and $\|P_n(\boldsymbol{u} * g - \boldsymbol{y})\|_p$ is a continuous function of g in the product topology and hence in the ρ-topology. Also note that computing this estimate does not require the knowledge of δ.

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

$$\|P_n(\boldsymbol{u} * \phi(P_n \boldsymbol{u}, P_n \boldsymbol{y}) - \boldsymbol{y})\|_p \le \|P_n(\boldsymbol{u} * h - \boldsymbol{y})\|_p = \|P_n \boldsymbol{d}\|_p \le \delta$$

and hence $\phi(P_n \boldsymbol{u}, P_n \boldsymbol{y}) \in S_n(\mathfrak{M}, \boldsymbol{u}, \boldsymbol{y}, \delta)$ 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 $h \in \mathfrak{M}$, define

$$A(h, \epsilon) \equiv \{g \in \mathfrak{M}: \rho(g, h) \ge D(\boldsymbol{u}, \mathfrak{M}, \delta) + \epsilon\}. \tag{3.8}$$

Also, consider the number

$$T(\epsilon) \equiv \sup_{h \in \mathfrak{M}} \sup_{g \in A(h, \epsilon)} T_{\boldsymbol{u}, \delta}(g, h) \tag{3.9}$$

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

$$B(\epsilon) = \{(g, h) \in \mathfrak{M}^2: \rho(g, h) \ge D(\boldsymbol{u}, \mathfrak{M}, \delta) + \epsilon\}.$$

It is clear that $B(\epsilon)$ is a closed set and hence compact in the ρ-topology, being a subset of $\mathfrak{M}^2$. Hence, $B(\epsilon)$ is also compact in the product topology. Now, $T_{\boldsymbol{u}, \delta}(g, h)$ is finite for all (g, h) in $B(\epsilon)$, by Lemma 3.7. Hence, by Lemma 3.8, $T(\epsilon)$ is finite.

Now if $n > T(\epsilon)$, then for any plant $h \in \mathfrak{M}$ and $\|\boldsymbol{d}\|_p \le \delta$, the estimate $\hat{h}^{(n)}$ generated by the algorithm must lie in the uncertainty set $S_n(\mathfrak{M}, \boldsymbol{u}, \boldsymbol{u} * h + d, \delta)$. Hence, by Lemma 3.6, $T_{\boldsymbol{u}, \delta}(\hat{h}^{(n)}, h) > n > T(\epsilon)$. This implies

$$\rho(\hat{h}^{(n)}, h) < D(\boldsymbol{u}, \mathfrak{M}, \delta) + \epsilon.$$

Since this holds for all h and $\boldsymbol{d}$, the convergence is indeed uniform. □

IV. Application of General Framework to Specific Problems

The above results state that under some compactness conditions on the model set, the optimal worst-case asymptotic error achievable by any identification algorithm is characterized by the function $D(\boldsymbol{u}, \mathfrak{M}, \delta)$, 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 $D(\boldsymbol{u}, \mathfrak{M}, \delta)$. In specific problems, we would like to find inputs $\boldsymbol{u}$ such that $D(\boldsymbol{u}, \mathfrak{M}, \delta)$ is small or, at the very least, vary continuously with the noise bound δ at $\delta = 0$. This would imply that identification accuracy is *robust* to measurement noise.

The value of the diameter of information $D(\boldsymbol{u}, \mathfrak{M}, \delta)$ is in general difficult to evaluate because it is the supremum over the diameter of all possible infinite-horizon uncertainty sets. However, if the ρ metric comes from a norm, it turns out that for an important class of model sets, $D(\boldsymbol{u}, \mathfrak{M}, \delta)$ 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 it follows from a basic result in information-based complexity theory [21].

Proposition 4.1: Suppose $\rho(g, h) \equiv \|g - h\|_{\mathfrak{X}}$ for some norm $\|\cdot\|_{\mathfrak{X}}$. If $\mathfrak{M}$ is a balanced convex subset of $\mathfrak{X}$, then the worst-case diameter is attained when the true plant and the disturbance are both 0. That is,

$$D(\boldsymbol{u}, \mathfrak{M}, \delta) \equiv \sup_{h \in \mathfrak{M}} \sup_{\|\boldsymbol{d}\|_\infty \le \delta} \operatorname{diam}(S_\infty(\mathfrak{M}, \boldsymbol{u}, \boldsymbol{u} * h + \boldsymbol{d}, \delta)) = \operatorname{diam}(S_\infty(\mathfrak{M}, \boldsymbol{u}, 0, \delta)).$$

Now we will apply the general results proved above to analyze specific identification problems. We take our input set $\mathfrak{U}$ to be $\overline{Bl}_\infty \equiv \{u: \|u\|_\infty \le 1\}$, where $\|u\|_\infty \equiv \sup_i |u_i|$. (The 1 is taken for normalization purpose.) The disturbance is assumed to be an l_∞ signal d, with $\|d\|_\infty \le \delta$.

A. Identification of Stable Plants in the l_1 Norm

Here the metric considered is $\rho(g, h) \equiv \|g - h\|_1$, and we restrict ourselves to stable plants with impulse responses of finite l_1 norm. We shall first prove a general lower bound for $D(\boldsymbol{u}, \mathfrak{M}, \delta)$ which holds for all inputs $\boldsymbol{u}$ and for a wide class of model sets.

Proposition 4.2: Assume the model set $\mathfrak{M}$ contains two plants at an l_1 distance of 2δ apart. Then for any number of experiments N and any set of inputs $\boldsymbol{u} \in \overline{Bl}_\infty^N$,

$$D(\boldsymbol{u}, \mathfrak{M}, \delta) \geq 2\delta.$$

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

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

$$\operatorname{diam}(S_\infty(\mathfrak{M}, \boldsymbol{u}, \boldsymbol{y}, \delta)) \geq \|g - h\|_1 = 2\delta.$$

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

We now demonstrate that in fact, for *all* balanced and convex model sets of stable plants, this lower bound can be reached using just *one* input, provided that it satisfies a persistent excitation property.

Definition 4.3: Let $\mathfrak{A}$ be the set of all finite sequences of 1's and -1's:

$$\mathfrak{A} \equiv \{(a_1, a_2, \cdots, a_k) : k \geq 1, a_i \in \{1, -1\}, \forall i\}. \quad (4.10)$$

The sequence $v \in \overline{Bl}_\infty$ is said to contain all finite sequences of 1's and -1's if for every finite sequence $a \in \mathfrak{A}$, there exist m, n such that $(v_m, v_{m+1}, \cdots, v_{m+n}) = a$.

Theorem 4.4: Assume $\mathfrak{M}$ is balanced and convex and contains only stable plants. If u^* contains all finite sequences of 1's and -1's, then

$$D(u^*, \mathfrak{M}, \delta) \leq 2\delta.$$

Proof: By Proposition 4.1, the diameter of information is given by the diameter of the uncertainty set centered at 0:

$$D(u^*, \mathfrak{M}, \delta) = \operatorname{diam}(S_\infty(\mathfrak{M}, u^*, 0, \delta)).$$

Consider any $g \in S_\infty(\mathfrak{M}, u^*, 0, \delta)$ and let $\epsilon > 0$. Since g is stable, there exists $\mathfrak{M}$ such that

$$\sum_{k=M+1}^{\infty} |g_k| < \epsilon. \quad (4.11)$$

Now consider the finite sequence

$$(\operatorname{sgn}(g_M), \operatorname{sgn}(g_{M-1}), \cdots, \operatorname{sgn}(g_0)) \in \mathfrak{A}$$

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

By definition of the sequence u^*, there exists m such that

$$u_m^* = \operatorname{sgn}(g_M), u_{m+1}^* = \operatorname{sgn}(g_{M-1}), \cdots,$$
$$u_{m+M}^* = \operatorname{sgn}(g_0).$$

We then have

$$\begin{aligned}
|(u^* * g)_{m+M}| &= \left| \sum_{k=0}^{m+M} u_{m+M-k}^* g_k \right| \\
&= \left| \sum_{k=0}^{M} u_{m+M-k}^* g_k + \sum_{k=M+1}^{m+M} u_{m+M-k}^* g_k \right| \\
&= \left| \sum_{k=0}^{M} \operatorname{sgn}(g_k) g_k + \sum_{k=M+1}^{m+M} u_{m+M-k}^* g_k \right| \\
&\geq \sum_{k=0}^{M} |g_k| - \sum_{k=M+1}^{m+M} |g_k| \\
&\geq \|g\|_1 - \epsilon. \qquad (4.12)
\end{aligned}$$

But $g \in S_\infty(\mathfrak{M}, u^*, 0, \delta)$, so $|(u^* * g)_{m+M}| \leq \delta$. Hence, it follows from inequality (4.12) that $\|g\|_1 \leq \delta + \epsilon$. Since this is true for every $\epsilon > 0$, it follows that $\|g\|_1 \leq \delta$ for any $g \in S_\infty(\mathfrak{M}, u^*, 0, \delta)$. Thus,

$$\begin{aligned}
D(\boldsymbol{u}, \mathfrak{M}, \delta) &= \operatorname{diam}(S_\infty(\mathfrak{M}, u^*, 0, \delta)) \\
&= \sup_{g \in S_\infty(\mathfrak{M}, u^*, 0, \delta)} 2\|g\|_1 \leq 2\delta. \qquad \square
\end{aligned}$$

An input satisfying the above condition has been proposed independently by Makila [25] for l_1 identification. 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.

Using the above result on the infinite-horizon diameter of information, we shall analyze the optimal asymptotic l_1 error for stable model sets.

The consistency result proved earlier applies to σ-compact model sets. The following technical lemma concerning the asymptotic l_1 error enables us to extend the result to model sets which are closure of σ-compact model sets as well.

Lemma 4.5: For any model set $\mathfrak{M}$, inputs $\boldsymbol{u} \in \overline{Bl}_\infty^N$, algorithm ϕ and $\delta \geq 0$,

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

where $\overline{\mathfrak{M}}$ is the closure of $\mathfrak{M}$ with respect to the l_1-topology on $\mathfrak{X}$. (The superscript "1" emphasizes that the metric used is the l_1 norm.)

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

$$\limsup_{n \to \infty} \|\phi(P_n \boldsymbol{u}, P_n(\boldsymbol{u} * h + \boldsymbol{d})) - h\|_1 \leq e_\infty^1(\phi, \boldsymbol{u}, \mathfrak{M}, x). \quad (4.13)$$

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

$$\limsup_{n\to\infty} \|\phi(P_n u, P_n(u * h + d)) - h\|_1 \tag{4.14}$$

$$\le \limsup_{n\to\infty} \|\phi(P_n u, P_n(u * h' + u*(h - h') + d)) - h'\|_1 + \epsilon. \tag{4.15}$$

Now, $\|u*(h - h') + d\|_\infty \le \delta + \epsilon$, so applying inequality (4.13) with $x = \delta + \epsilon$,

$$\limsup_{n\to\infty} \|\phi(P_n u, P_n(u * h' + u*(h - h') + d)) - h'\|_1 \le e^1_\infty(\phi, \mathfrak{M}, u, \delta + \epsilon). \tag{4.16}$$

It follows that

$$\limsup_{n\to\infty} \|\phi(P_n u, P_n(u * h + d)) - h\|_1 \le e^1_\infty(\phi, \mathfrak{M}, u, \delta + \epsilon) + \epsilon. \tag{4.17}$$

Letting ϵ go to 0 gives the desired result. □

We now show that we can get very good asymptotic error even if there is no additional prior knowledge about the plant other than the fact that it is stable.

Proposition 4.6: Take the model set to be l_1, the space of all stable plants. There is a single experiment, using any input $u^* \in \overline{Bl}_\infty$ containing all sequences of 1's and -1's, such that for every $\delta \ge 0$, the optimal asymptotic l_1 error satisfies

$$E^1_\infty(u^*, l_1, \delta) \le 2\delta.$$

Proof: The space l_1 is separable, i.e., it is a closure of a countable set $\mathfrak{M}_\infty$. Since a countable set is clearly σ-compact, by Theorem 3.4, there is an algorithm ϕ^* such that for every $\delta \ge 0$ and inputs u,

$$e^1_\infty(\phi^*, \mathfrak{M}_\infty, u, \delta) \le D(u, \mathfrak{M}_\infty, \delta). \tag{4.18}$$

Now, using any input u^* containing all sequences of 1's and -1's, we have

$$\begin{aligned} e^1_\infty(\phi^*, \mathfrak{M}_{\text{stab}}, u^*, \delta) &\le \lim_{x\downarrow\delta} e^1_\infty(\phi^*, \mathfrak{M}_\infty, u^*, x) && \text{by Proposition 4.5} \\ &\le \lim_{x\downarrow\delta} D(u^*, \mathfrak{M}_\infty, x) \\ &\le 2\delta, && \text{by Theorem 4.4.} \end{aligned}$$ □

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. One should also note that there are many other choices of decomposing the model set into compact sets. The decomposition should be done to facilitate a more efficient implementation of the identification algorithm. We will discuss this at the end of this section.

Next, we look at the issue of uniform convergence. For the model set l_1, it can at once be seen that although convergence to a small asymptotic error is possible, such convergence cannot be uniform.

Proposition 4.7: Let ϕ be any algorithm and u be any input. Then for every n and for every M, there exists an $h \in \mathfrak{M}_{stab}$ such that

$$\|\phi(P_n u, P_n(u * h)) - h\|_1 > M.$$

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. □

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

Proposition 4.8: Let $\mathfrak{M} \subset \mathfrak{M}_{\text{stab}}$ be a compact set (in the l_1-topology) or a subset of a compact set in 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 $h \in \mathfrak{M}$ and all $\|d\|_\infty \le \delta$, 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.

Common examples of such compact model sets are the uniformly stable ones, of the form $M_s(g) \equiv \{h: |h_i| \le |g_i| \text{ for all } i\}$ where g is any stable plant. The specific model sets considered in [8] and [9] belong to this class.

Identification Algorithms for Stable Plants: For certain parameterizations of the space of stable plants, it is possible to device algorithms based on the *Occam's Razor Principle* that involve linear programming problems. Define the compact sets:

$$\mathfrak{M}_k = \{h \in l_1 | : |h_i| \le kM, h_i = 0 \ \forall i \ge k\}$$

and M is any positive real number. It can be immediately seen that

$$l_1 = \text{closure of } \bigcup_{k=1}^{\infty} \mathfrak{M}_k.$$

Fix some tolerance level ϵ. The estimator can be described as picking a feasible element in the set

$$\mathfrak{M}_k \cap S_n(\mathfrak{M}, u, y, \delta + \epsilon)$$

for any input–output pair. Of course, this set is characterized by linear constraints and finding a feasible plant is equivalent to solving a linear programming problem. The estimate is picked from the smallest $\mathfrak{M}_k$ for which the above set is not empty.

Suppose that the model set is equal to $\mathfrak{M}_s(g)$ where $g \in l_1$ and $g_i = 0 \ \forall i \ge l$. This set contains only FIR plants of length l, with a bound on the impulse response. For this model set the near-optimal algorithm ϕ^* is given by

$$\phi^*(P_n u, P_n y) = \underset{|h_i| \le |g_i|,\, i = 0, 1, \cdots l}{\arg\min} \|P_n(y - u * h)\|_\infty$$

which is computable by linear programming. We finally note that work on algorithms is still an active area of research [34].

B. H_∞ Identification of Stable Rational Plants

We now analyze optimal identification using the model set RH_∞, space of all stable plants with rational transfer functions. The error metric used is the H_∞ norm. The model set RH_∞ is σ-compact in the H_∞-topology. (For example, it can be decomposed as a countable union of compact sets of the form $\{h: |h_n| \le A\alpha^n\}$ with A tending to infinity and α tending to 1.) Convergence in H_∞ implies component-wise convergence of the impulse response in each of these sets. Hence, the consistency result applies and we are reduced to the analysis of the infinite-horizon diameter of information.

Since the H_∞-norm of a plant is always upper bounded by its l_1 norm, Theorem 4.4 implies that, measured in the H_∞ norm, the infinite horizon diameter of information $D(u^*, RH_\infty, \delta)$ using an input u^* containing all finite sequences of 1's and -1's is also bounded by 2δ. Hence, the worst-case asymptotic error using this input is also bounded by 2δ. The following result shows that this input is optimal to within a factor of two.

Proposition 4.9: For any number of experiments N and any choice of inputs $\boldsymbol{u} \in \overline{Bl}_\infty^N$, the H_∞ infinite-horizon diameter of information satisfies:

$$D(\boldsymbol{u}, RH_\infty, \delta) \ge 2\delta.$$

Proof: The proof is trivial. Take $g = (\delta, 0, 0, \cdots)$, $h = (-\delta, 0, 0, \cdots)$, $\boldsymbol{d} = -\delta\boldsymbol{u}$, $\boldsymbol{d}' = \delta\boldsymbol{u}$. Then $\boldsymbol{u} * g + \boldsymbol{d} = \boldsymbol{u} * h + \boldsymbol{d}'$ so $D(\boldsymbol{u}, RH_\infty, \delta) \ge \|g - h\|_{H_\infty} = 2\delta$. □

A similar result on frequency response experiments is given by [9].

C. Identification of Unstable Plants in the Gap Metric

Our general framework of optimal asymptotic identification applies, to a large extent, to unstable as well as stable systems. In particular, the consistency and uniform convergence results, for arbitrary inputs, hold regardless of whether the model set contains stable or unstable systems. There is, however, an important issue in the identification of unstable systems which is not dealt with in this framework. While stable systems can be identified in the open-loop, identification experiments for unstable systems are almost always performed in the closed-loop to avoid unbounded outputs. As opposed to open-loop identification, there is no complete freedom in choosing the inputs $\boldsymbol{u}$ for closed-loop identification experiments, as there is a coupling between the input and the output. This makes the experiment design problem much more difficult. In this section, we shall ignore the coupling and confine ourselves to deriving necessary and sufficient conditions on the inputs for accurate asymptotic identification of unstable systems. The question of whether one can design closed-loop experiments to achieve such conditions is left open.

An appropriate error metric to use for unstable plants is the *gap metric* [6], [33], [42]. The important property of the gap metric is that it generates the graph topology [40], which is the weakest topology in which closed-loop stability is a robust property, or in which the closed-loop system varies continuously as a function of the open-loop system. Intuitively, this means that identifying plants accurately in the gap metric is the least that one must do to be able to design controllers to guarantee that the closed-loop performance will be close to the desired.

The gap between two possibly unstable plants is given in terms of their *graphs*, so we will first define this notion. The graph G_h of a plant h is a subset of the space $l_2 \times l_2$, defined by

$$G_h \equiv \{(x, h * x): x \in l_2, h * x \in l_2\}.$$

Thus, the graph of a plant describes its behavior on bounded-energy inputs which yield bounded-energy outputs. The directed gap between two graphs G_h and G_g is defined as

$$\vec{\delta}(G_h, G_g) \equiv \sup_{x \in G_h, \|x\|_2 \le 1} \inf_{y \in G_g} \|x - y\|_2.$$

The gap between two plants is given by the maximum of the two directed gaps between the two graphs:

$$\delta(g, h) \equiv \max\left(\vec{\delta}(G_g, G_h), \vec{\delta}(G_h, G_g)\right).$$

It can be verified that the gap is indeed a metric, and that its value is always bounded between 0 and 1.

In the analysis below, we shall restrict ourselves to the space of finite-dimensional systems, $\mathfrak{M}_{fd}$, with rational z-transform.[1] In this space, convergence in the graph topology can be expressed in terms of the coprime factors: $P_i \to P$ in the graph topology iff there exist co-prime factorizations $P_i = N_i/D_i$, $P = N/D$ such that $N_i \to N$ and $D_i \to D$ in the H_∞-topology. Results obtained for finite-dimensional plants are also valid for infinite-dimensional systems that can be approximated by finite-dimensional systems in the gap metric.

To apply the consistency results we proved earlier, we have to investigate the topological properties of M_{fd}.

Proposition 4.10: Let p, q be nonnegative integers, k, α be positive real numbers and $\mathfrak{M}_{fd}(p, q, K, \alpha)$ be the class of all finite-dimensional systems having z-transforms

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

with bounded parameters: $|a_i| \le K$ and $|b_i| \le k$ for all i, and with the distance between any pole-zero pair $\ge \alpha$. $\mathfrak{M}_{fd}(p, q, K, \alpha)$ is compact in the graph topology, and on this set the graph topology is finer than the product topology.

Proof: Let $\{P_i(z)\}$ be a sequence of plants in $\mathfrak{M}_{fd}(p, q, K, \alpha)$, and suppose $P_i = N_i/D_i$, with deg $N_i \le p$, deg $D_i = q$, D_i monic, and the coefficients of N_i and D_i bounded by K. Clearly, N_i and D_i lie in sets which are compact in the H_∞-topology. Hence, there exist a subsequence $N_{k_i} \to N^*$ and $D_{k_i} \to D^*$. We now verify that $P^* \equiv N^*/D^*$ is in $\mathfrak{M}_{fd}(p, q, K, \alpha)$. We first note that H_∞

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

convergence of polynomials of bounded degree is equivalent to convergence of their coefficients. Hence, deg $N^* \le p$, deg $D^* = q$, D^* is monic, and their coefficients are bounded by K. Moreover, since the location of the zeros of a polynomial is continuous of its coefficients, the zeros of N_{k_i}, D_{k_i} must converge to those of N^*, D^*, respectively, and the separation between poles and zeros is maintained at a distance of at least α. Hence, $P^* \in \mathfrak{M}_{fd}(p, q, K, \alpha)$, and $P_{k_i} \to P^*$ in the graph topology. This shows that $\mathfrak{M}_{fd}(p, q, K, \alpha)$ is compact in the graph topology. Also, in $\mathfrak{M}_{fd}(p, q, K, \alpha)$, convergence in the graph implies convergence in the coefficients of the rational transfer function, which in turn implies the convergence in each component of the impulse response. This latter fact follows by inspection of the inversion formula for z-transforms. □

It is clear that the space of all finite-dimensional systems $\mathfrak{M}_{fd}$ is a countable union of sets of the form $\mathfrak{M}_{fd}(p,q, K, \alpha)$. It then follows that Theorem 3.4 can be applied on M_{fd} equipped with the gap metric, and the infinite-horizon diameter of information $D_{\text{gap}}(\boldsymbol{u}, \mathfrak{M}_{fd}, \delta)$ characterizes the optimal asymptotic error $\bar{E}_{\infty}(\boldsymbol{u}, \mathfrak{M}_{fd}, \delta)$.

We shall first derive necessary conditions on the inputs $\boldsymbol{u}$ for the robustness of the asymptotic error to measurement noise, i.e., when $D_{\text{gap}}(\boldsymbol{u}, \mathfrak{M}_{fd}, \delta)$ approaches 0 as δ approaches 0. This is in terms of the notion of *stability testing*: inputs $\boldsymbol{u} \equiv [u^{(1)}, u^{(2)}, \cdots, u^{(N)}]$ are said to be able to *test the stability* of plants if for every unstable $h \in \mathfrak{M}_{fd}$ at least one of the inputs $u^{(i)}$ yields an unbounded output. We have the following result on the loss of robustness when the inputs are not rich enough to test stability.

Proposition 4.11: If the inputs $\boldsymbol{u}$ cannot test stability, then $D_{\text{gap}}(\boldsymbol{u}, \mathfrak{M}_{fd}, \delta) = 1$ for all $\delta > 0$.

Proof: Let $\delta > 0$. Consider the infinite-horizon uncertainty set centered at the origin:

$$S_{\infty}(\mathfrak{M}_{fd}, \boldsymbol{u}, 0, \delta) = \{g \in \mathfrak{M}_{fd} : \|\boldsymbol{u} * g\|_{\infty} \le \delta\}.$$

Since $\boldsymbol{u}$ cannot test the stability of plants in M_{fd}, there must be an unstable plant $h \in \mathfrak{M}_{fd}$ such that $\boldsymbol{u} * h$ is bounded; by appropriate scaling, we can assume that $h \in S_{\infty}(\mathfrak{M}_{fd}, \boldsymbol{u}, 0, \delta)$. Since the zero plant is also in this uncertainty set and the gap distance between the zero plant and any unstable plant is 1 [6], the diameter of this uncertainty set must be 1. Hence, the diameter of information, which is the diameter of the largest uncertainty set, is also 1. □

We now give explicit necessary and sufficient conditions for inputs to be able to test stability. We begin with two definitions.

Definition 4.12: 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 4.13: A sequence u is said to excite at frequency $\omega \in [0, 2\pi]$ if

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

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

We shall now give the following result, the proof of which can be found in the Appendix.

Theorem 4.14: $\mathfrak{M}_{fd}$ is testable for stability by bounded inputs $u^{(1)}, \cdots, u^{(N)}$ if and only if the inputs have the following properties:

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

2) $$\bigcap_{i=1}^{N} Z(u^{(i)}) = \varnothing.$$

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

We have the following corollary.

Corollary 4.15: $\mathfrak{M}_{fd}$ 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 nonexistence of a bounded input having both the properties required by Corollary 4.15 has been established. However, bounded inputs which excite at all frequencies do exist. In fact, Lusin [19] has constructed a sequence which excites at all frequencies despite the fact that the sequence actually tends to 0.

Stability testing is a necessary property the inputs must satisfy in order to have robustness in the asymptotic error. It will now be shown that stability testing combined with the property of containing all finite sequences of 1's and -1's are in fact sufficient to guarantee robustness.

Theorem 4.16: If the inputs $\boldsymbol{u}$ can test stability and at least one of them contains all finite sequences of 1's and -1's then for all $\delta \ge 0$,

$$D_{\text{gap}}(\boldsymbol{u}, \mathfrak{M}_{fd}, \delta) \le 2\delta.$$

Proof: Consider now the infinite-horizon uncertainty set $S_{\infty}(\mathfrak{M}_{fd}, \boldsymbol{u}, 0, \delta)$ centered at the origin. Since all the plants in this set give zero output on the inputs and the inputs test stability, all the plants in this set must be stable. Moreover, one of the inputs contains all finite sequences of 1's and -1's. We are now in a similar situation as in Theorem 4.4, which applies to the stable plant case. Exact arguments as in the proof of that theorem show that the diameter of this uncertainty set *measured in the* l_1 *norm* is bounded by 2δ. Since $\mathfrak{M}_{fd}$ is balanced and convex, the diameter of information equals diameter of this set (measured in the l_1 norm). Finally, by a result proved in the Appendix, the gap distance between two plants is always bounded by the H_{∞} distance, and therefore also by the l_1 distance. Hence, the diameter of information $D_{gap}(\boldsymbol{u}, \mathfrak{M}_{fd}, \delta)$ measured in the gap metric is bounded by the diameter of information measured in the l_1 norm, and hence also bounded by 2δ. □

We will now exhibit two inputs which have the above desired properties. First, it will be demonstrated that any input that contains all finite sequences of 1's and -1's excites at all frequencies.

Proposition 4.17: 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 $\Sigma_k |\cos k\omega_0|$ is divergent, so we can find an integer L such that $\Sigma_{k=0}^{L} |\cos k\omega_0| > M$. By the definition of the sequence u, there exists an integer n_1 such that

$$\begin{aligned}&(u_{n_1}, u_{n_1+1}, \cdots, u_{n_1+L})\\ &= (1, \operatorname{sgn}(\cos\omega_0), \operatorname{sgn}(\cos 2\omega_0), \cdots, \operatorname{sgn}(\cos L\omega_0)).\end{aligned} \tag{4.19}$$

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 $\limsup_{n\to\infty} |\Sigma_{k=0}^{n} u_k e^{-jk\omega_0}| = \infty$. □

Using two inputs, one of which contains all finite sequences of 1's and -1's and the another the unit impulse, will suffice to test stability, since the former excites at all frequencies and the latter's z-transform has no zeros in the unit disk. It follows immediately from the Theorem 4.16 that an optimal worst-case gap error of 2δ can be achieved with these two inputs.

This result shows that for finite-dimensional plants, identification in the gap metric can be performed *robust* to the noise level δ, i.e., as δ goes to zero, the identification error also goes to zero. However, we have not yet shown that the two experiments are optimal or near optimal. A lower bound to the optimal asymptotic gap error using *any* bounded inputs will now be derived. This will show that for small δ, the above experiment design is no more than a factor of two from optimality.

Proposition 4.18: For any N and inputs $\boldsymbol{u} \in \overline{Bl}_\infty^N$, the optimal worst-case asymptotic gap error for finite-dimensional plants satisfies

$$E_\infty^g(\boldsymbol{u}, \mathfrak{M}_{fd}, \delta) \geq \frac{\delta}{\sqrt{1+\delta^2}}.$$

Proof: To prove this result, it suffices to show that the infinite-horizon gap diameter of information satisfies

$$D_{\text{gap}}(\boldsymbol{u}, \mathfrak{M}_{fd}, \delta) \geq 2\frac{\delta}{\sqrt{1+\delta^2}}.$$

We make use of the following lower bound for the gap metric [44]:

$$\delta(h, 0) \geq \frac{\|h\|_{H_\infty}}{\sqrt{1+\|h\|_{H_\infty}^2}}.$$

Now,

$$\begin{aligned}&D_{\text{gap}}(\boldsymbol{u}, \mathfrak{M}_{fd}, \delta)\\ &= \sup_{h\in\mathfrak{M}_{fd}} \sup_{\|\boldsymbol{d}\|_\infty\leq\delta} \operatorname{diam}_{\text{gap}} S_\infty(\mathfrak{M}_{fd}, \boldsymbol{u}, h * \boldsymbol{u} + \boldsymbol{d}, \delta)\\ &\geq \operatorname{diam}_{\text{gap}} S_\infty(\mathfrak{M}_{fd}, \boldsymbol{u}, 0, \delta)\\ &= \sup_{g\in\mathfrak{M}_{fd},\, \|g*\boldsymbol{u}\|_\infty\leq\delta} 2\delta(g, 0)\\ &\quad \text{since } \delta(g,0) = \delta(-g,0)\\ &\geq \sup_{g\in\mathfrak{M}_{fd},\, \|g*\boldsymbol{u}\|_\infty\leq\delta} 2\frac{\|g\|_{H_\infty}}{\sqrt{1+\|g\|_{H_\infty}^2}}\\ &\quad \text{using the lower bound to the gap}\\ &\geq \frac{2\delta}{\sqrt{1+\delta^2}}\\ &\quad \text{choosing } g \text{ to be an impulse with magnitude } \delta.\end{aligned}$$

□

Finally, we note that this theorem has interesting implications to identification in the closed loop. To accurately estimate the plant, it is necessary that the input satisfies the conditions in Theorem 4.14. In general it is not known whether there exists one input with that property. If not, then more information about the model set should be known. An example of such information is the knowledge of a stabilizing controller of the plant to be identified. Details on this can be found in [29], [39].

V. 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 to within a factor of two of this bound asymptotically. These results hold for any error metric and disturbance norm. Second, for specific identification problems, characterization of inputs which makes this infinite-horizon diameter of information small is given. In particular, we considered identification in both the l_1 and the H_∞ norms for stable plants, and in the gap metric for unstable finite-dimensional plants of arbitrary order. The significance of these error metrics is that if the worst-case error is small in these metrics, methods exist for synthesizing controllers to achieve robust performance [2], [4].

The results show that accurate identification is possible in the worst case for a specific choice of inputs depending on the model set. For identification in the l_1 norm, algorithms for computing estimates are based on linear programming and are easily implementable. For the identification in the gap metric, robust identification was shown to be more or less equivalent to stability testing. This has important implications on closed-loop identification in which one does not have direct access to the input.

There are many issues in worst-case identification that need to be resolved. The issue of computational complexity and implementation of the algorithm is a central issue. In particular, it is beneficial to relate the complexity of the model set to the complexity of the required experiments and the algorithms. Another issue is the relationship between the identification in the frequency domain and the time domain, particularly as it relates to algorithms and complexity. Deeper study of identification of unstable plants in a closed-loop setting is needed. The relations of all of this to adaptive control is of course one of the prime motivations for this work and will be the subject of future research.

VI. Acknowledgments

We would like to thank Sanjeev Kulkarni, Sanjoy Mitter, and George Zames for very helpful discussions.

Appendix

A. Proof of Theorem 4.14

To prove this result, we need the following lemma, the proof of which is elementary but tedious, and can be found in [38].

Lemma A.1: Let $u \in \overline{Bl}_\infty$ 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 4.14.

Proof:

(if part)

Let $u^{(1)}, u^{(2)}, \cdots, u^{(N)} \in \overline{Bl}_\infty$ be N inputs satisfying properties (1) and (2). Let $h \in \mathfrak{M}_{fd}$ with a rational z-transform $H(z)$, and assume that the outputs $u^{(i)} * h$, $i = 1, \cdots, N$, 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

$$H(z) = H_u(z) + H_s(z) \tag{A.20}$$

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 $\omega_1, \omega_2, \cdots, \omega_M$. Then $H_u(z)$ can be decomposed as

$$H_u(z) = \sum_{i=1}^{M} H_i(z) \tag{A.21}$$

where

$$H_i(z) = \frac{\sum_{k=0}^{L_i-1} \alpha_{ik} z^k}{(z - e^{j\omega_i})^{L_i}} \tag{A.22}$$

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 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 $\bigcup_{i=1}^{N} \Omega(u^{(i)}) = [0, 2\pi]$ the frequency ω_1 lies in $\Omega(v)$ for some input $v \in \{u^{(1)}, \cdots, u^{(N)}\}$. By Proposition A.1,

$$y^{(1)} = v * h^{(1)} \notin l_\infty \tag{A.23}$$

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

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

$$x_{n+1}^{(1)} = A_1 x_n^{(1)} + B_1 v_n, \qquad y_n^{(1)} = C_1 x_n^{(1)} \tag{A.24}$$

for some matrices A_1, B_1, C_1.

Since $y^{(1)}$ is unbounded but v is bounded, it follows from (A.24) that the modal states $x^{(1)}$ must be unbounded given input v. But the overall state 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

$$x_{n+1} = Ax_n + Bv_n, \qquad y_n = Cx_n. \tag{A.25}$$

From (A.25), a sequence of equations is obtained as

$$y_n = Cx_n$$
$$y_{n+1} = CAx_n + CBv_n$$
$$\vdots$$
$$y_{n+L-1} = CA^{L-1}x_n + \sum_{i=0}^{L-2} CA^iBv_n.$$

Let

$$\mathbf{y}_n = \begin{bmatrix} y_n \\ y_{n+1} \\ \vdots \\ y_{n+L-1} \end{bmatrix}, \qquad Q_0(A,C) = \begin{bmatrix} C \\ CA \\ \vdots \\ CA^{L-1} \end{bmatrix},$$

$$E = \begin{bmatrix} 0 \\ CB \\ \vdots \\ \sum_{i=0}^{L-2} CA^iB \end{bmatrix}.$$

The sequence of output equations can then be written as

$$\mathbf{y}_n = Q_0(A,C)x_n + Ev_n. \qquad \text{(A.26)}$$

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 x_n becomes unbounded and v_n is bounded, the output $\mathbf{y}_n$ 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 $u^{(1)},\cdots,u^{(N)}$ can test stability in $\mathfrak{M}_{fd}$.

(only-if part)

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

Suppose the first condition is not satisfied; consider an $\omega_0 \in [0,2\pi]$ but $\omega_0 \notin \bigcup_{i=1}^N \Omega(u^{(i)})$. Consider the unstable system $h_n = \cos(n\omega_0)$. Lemma A.1(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 $\mathfrak{M}_{fd}$. 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,

$$\sum_{k=0}^{\infty} u_k^{(i)} r_0^k e^{jk\omega_0} = 0, \qquad \forall i. \qquad \text{(A.27)}$$

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

$$\sum_{k=0}^{\infty} u_k^{(i)} r_0^k e^{-jk\omega_0} = 0 \qquad \forall i. \qquad \text{(A.28)}$$

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

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

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

B. An Inequality Between the Gap and H_∞ Distances

Proposition B.1: Let h and g be two plants. Then

$$\delta(g,h) \le \|h - g\|_{H_\infty}.$$

Proof: We assume that $\|h - g\|_{H_\infty} < \infty$ otherwise there is nothing to prove. Now,

$$\delta(g,h) = \max\left(\vec{\delta}(G_g, G_h), \vec{\delta}(G_h, G_g)\right)$$

where

$$G_h \equiv \{(u, h * u) \in l_2 : x \in l_2, h * x \in l_2\}$$

and

$$\vec{\delta}(G_h, G_g) \equiv \sup_{x \in G_h, \|x\|_2 \le 1} \inf_{y \in G_g} \|x - y\|_2.$$

Now, since $\|g - h\|_{H_\infty} < \infty$

$$(u, h * u) \in G_h \Leftrightarrow (u, g * u) \in G_g.$$

We have

$$\begin{aligned}\vec{\delta}(G_h, G_g) &\le \sup_{h * u \in l_2, \|u\|_2 \le 1} \inf_{y \in G_g} \|(u, h * u) - y\|_2 \\ &\le \sup_{h * u \in l_2, \|u\|_2 \le 1} \|(u, h * u) - (u, g * u)\|_2 \\ &= \sup_{h * u \in l_2, \|u\|_2 \le 1} \|(h - g) * u\|_2 \\ &\le \|h - g\|_{H_\infty}.\end{aligned}$$

Hence, the result follows. □

References

[1] A. Blumer, A. Ehrenfeucht, D. Haussler, and M. Warmuth, "Occam's razor," *Inf. Proc. Lett.*, vol. 24, pp. 377–380, 1987.

[2] M. A. Dahleh and M. H. Khammash, "Controller design in the presence of structured uncertainty," to appear in *Automatica* special issue on robust control.

[3] M. A. Dahleh, T. Theodosopoulos, and J. N. Tsitsiklis, "The sample complexity of worst-case identification of F. I. R. linear systems," *Syst. Contr. Lett.*, to be published.

[4] J. C. Doyle, "Analysis of feedback systems with structured uncertainty," *IEE Proc.*, vol. 129, pp. 242–250, 1982.

[5] 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.

[6] T. T. Georgiou and M. C. Smith, "Optimal robustness in the gap metric," *IEEE Trans. Automat. Contr.*, vol. 35, no. 6, pp. 673–686, 1990.

[7] G. C. Goodwin, "Experiment design in system identification" in *Encyclopedia of Systems and Control*, M. Singh, Ed. New York: Pergamon Press, 1987.

[8] G. Gu and P. P. Khargonekar, "Linear and nonlinear algorithms for identification in $\mathfrak{H}_\infty$ with error bounds, *IEEE Trans. Automat. Contr.*, vol. 37, no. 7, July 1992.

[9] A. J. Helmicki, C. A. Jacobson, and C. N. Nett, "Identification in H_∞: A robust convergent nonlinear algorithm," in *Proc. 1989 Int. Symp. Math. Theory Networks Syst.*, 1989.

[10] ——, "Identification in H_∞: Linear algorithms," in *Proc. 1990 Amer. Contr. Contr.*, pp. 2418–2423.

[11] A. J. Helmicki, C. A. Jacobson, and C. N. Nett, "Control-oriented system identification: A worst-case deterministic approach in H_∞," *IEEE Trans. Automat. Contr.*, vol. 36, no. 10, Oct. 1991.

[12] C. A. Jacobson and C. N. Nett, "Worst-case system identification in l_1: Optimal algorithms and error bounds," in *Proc. 1991 Amer. Contr. Conf.*, June 1991.

[13] B. Kacewicz and M. Milanese, "On the optimal experiment design in the worst-case l_1 *system identification*," submitted to CDC.

[14] B. Kacewicz and L. Plaskota, "Noisy information for linear systems in asymptotic setting," *J. Complexity*, vol. 7, pp. 35–57, 1991.

[15] J. M. Krause, G. Stein, and P. P. Khargonekar, "Robust performance of adaptive controllers with general uncertainty structure," in *Proc. 29th Conf. Decision Contr.*, 1990, pp. 3168–3175.

[16] 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.

[17] M. K. Lau, R. L. Kosut, and S. Boyd, "Parameter set estimation of systems with uncertain nonparametric dynamics and disturbances," in *Proc. 29th Conf. Decision Contr.*, 1990, pp. 3162–3167.

[18] L. Lin, L. Wang, and G. Zames, "Uncertainty principle and identification n-widths for LTI and slowly varying systems," ACC, Chicago, IL, 1992.

[19] A. Markushevich, *Theory of Functions of A Complex Variable*. Englewood Cliffs, NJ: Prentice-Hall, vol. I, 1965.

[20] R. K. Mehra, "Optimal input signals for parameter estimation in dynamic systems—A survey and new results," *IEEE Trans. Automat. Contr.*, vol. AC-19, pp. 753–768, 1974.

[21] 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. New York: Plenum, 1977.

[22] 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. Automat. Contr.*, vol. AC-27, pp. 408–414, 1982.

[23] M. Milanese and R. Tempo, "Optimal algorithm theory for robust estimation and prediction," *IEEE Trans. Automat. Contr.*, vol. AC-30, pp. 730–738, 1985.

[24] M. Milanese, "Estimation theory and prediction in the presence of unknown and bounded uncertainty: A survey," in *Robustness in Identification and Control*, M. Milanese, R. Tempo, A. Vicino, Eds. New York: Plenum Press, 1989.

[25] P. M. Makila, "Robust identification and galois sequences," *Int. J. Contr.*, vol. 54, pp. 1189–1200.

[26] P. M. Makila and J. R. Partington, "Robust approximation and identification in H_∞," in *Proc. 1991 Amer. Cont. Conf.*, June, 1991.

[27] J. Munkres, *Topology—A First Course*. Englewood Cliffs, NJ: Prentice-Hall, 1975.

[28] J. P. Norton, "Identification and application of bounded-parameter models," *Automatica*, vol. 23, no. 4, pp. 497–507, 1987.

[29] J. R. Partington and P. M. Makila, "Worst-case identification of stabilizable systems," submitted.

[30] J. Pearl, "On the connection between the complexity and credibility of inferred models," *Int. J. General Syst.*, vol. 4, pp. 255–264, 1978.

[31] K. Poolla and A. Tikku, "On the time complexity of worst-case system identification," preprint, 1992.

[32] L. Pronzato and E. Walter, "Experiment design in bounded-error context: Comparison with D-optimality," *Automatica*, vol. 25, no. 3, pp. 383–391, 1989.

[33] A. K. El-Sakkary, "The Gap Metric: Robustness of Stabilization of Feedback Systems," *IEEE Trans. Automat. Contr.*, vol. AC-30, pp. 240–247, 1985.

[34] T. Theodosopoulos, "Worst-case identification in l_1: Algorithms and complexity," thesis in preparation.

[35] J. F. Traub and H. Wozniakowski, *A General Theory of Optimal Algorithms*. New York: Academic, 1980.

[36] J. F. Traub, G. Wasilkowski, and H. Wazniakowski, *Information-Based Complexity*. New York: Academic, 1988.

[37] D. N. C. Tse, M. A. Dahleh, and J. N. Tsitsiklis, "Optimal and Robust Identification in the l_1 norm," in *Proc. 1991 Amer. Contr. Conf.*, June 1991.

[38] D. N. C. Tse, "Optimal and robust identification under bounded disturbances," master's thesis, Dept. Elec. Eng. Comp. Sci., M.I.T., Feb. 1991.

[39] D. N. C. Tse, M. A. Dahleh, and J. N. Tsitsiklis, "Optimal asymptotic worst-case identification with applications on l_1 and gap metrics," in *Proc. Int. Symp. MTNS-91*, pp. 325–330.

[40] M. Vidyasagar, *Control System Synthesis*. Cambridge, MA: M.I.T. Press, 1985.

[41] G. Zames, "On the metric complexity of casual linear systems: ϵ-entropy and ϵ-dimension for continuous-time," *IEEE Trans. Automat. Contr.*, vol. 24, Apr. 1979.

[42] G. Zames and A. El-Sakkary, "Unstable systems and the gap metric," in *Proc. Allerton Conf.*, pp. 380–385, Oct. 1980.

[43] M. Zarrop, *Optimal Experimental Design for Dynamic System Identification*. New York: Springer-Verlag, 1979.

[44] S. Zhu, M. Hautus, and C. Praagman, "A lower and upper bound for the gap metric," in *Proc. 28th Conf. Decision Contr.*, Dec. 1989, pp. 2337–2341.

David N. C. Tse (S'91) was born in Hong Kong, in 1966. He received the B.A.Sc. degree in systems design engineering from University of Waterloo, Canada, in 1989, and the M.S. degree in electrical engineering from Massachusetts Institute of Technology, Cambridge, MA, in 1991.

He is currently a Ph.D. student in the Laboratory for Information and Decision Systems at MIT. His research interests include data compression, stochastic processes, system identification, and information theory.

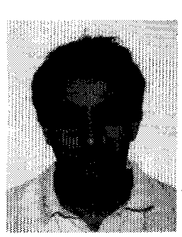

Munther A. Dahleh (S'84–M'87) was born in August 1962. He received the B.S.E.E. degree from Texas A & M University, College Station, TX in 1983, and the Ph.D. degree from Rice University, Houston, TX 1987.

Since 1987, he has been with the Department of Electrical Engineering and Computer Science, Massachusetts Institute of Technology, Cambridge, MA, where he is currently an Associate Professor. In July 1992 he was awarded the Finmeccanica career development chair. He has held consulting positions with NASA and the C.S. Draper Laboratory since 1988, and with FIAT Research Center since 1991. His current interests include robust control, identification of uncertain systems and adaptation, and control of time-varying and infinite-dimensional systems.

Dr. Dahleh is currently serving as an Associate Editor for the *Systems and Control Letters*.

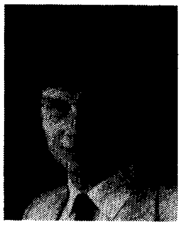

John N. Tsitsiklis (S'80–S'81–M'81–M'83) was born in Thessaloniki, Greece, in 1958. He received the B.S. degree in mathematics, in 1980, and the B.S., M.S., and Ph.D. degrees in electrical engineering, in 1980, 1981, and 1984, respectively, all from the Massachusetts Institute of Technology, Cambridge, MA.

During the academic year 1983–1984, he was an acting Assistant Professor of Electrical Engineering at Stanford University, Stanford, CA. Since 1984, he has been with the Department of Electrical Engineering and Computer Science at the Massachusetts Institute of Technology, where he is currently an Associate Professor. His research interests are in the areas of systems and control theory, parallel and distributed computation, and operations research. He has coauthored more than 50 papers in these areas. He is also a coauthor (with D. Bertsekas) of *Parallel and Distributed Computation: Numerical Methods*, 1989. He has been a recipient of an IBM Faculty Development Award, 1983, an NSF Presidential Young Investigator Award, 1986, an Outstanding Paper Award by the IEEE Control Systems Society (for a paper coauthored with M. Athans, 1986), and of the Edgerton Faculty Achievement Award by M.I.T., 1989. He has been a plenary speaker at the *1992 IEEE Conference on Decision and Control*. Finally, he is an Associate Editor of *Automatica* and *Applied Mathematics Letters*, and has been an Associate Editor of the IEEE TRANSACTIONS ON AUTOMATIC CONTROL.