A Tighter Bound for the Echo State Property

Michael Buehner and Peter Young

Abstract

This letter provides a brief explanation of echo state networks and provides a rigorous bound for guaranteeing asymptotic stability of these networks. The stability bounds presented here could aid in the design of echo state networks that would be applicable to control applications where stability is required.

Index Terms

echo state networks, weighted operator norms, recurrent neural networks, nonlinear systems, Lyapunov stability, robust controls

I. INTRODUCTION

Artificial neural networks may be used in control system applications such as modelling nonlinear system dynamics and control of nonlinear systems. Two types of networks are used for these applications, namely feed-forward neural networks (FFNN) and recurrent neural networks (RNN). FFNN are attractive since they are easy to train in a stable manner (e.g. using back propagation), but are limited in the sense that they are only capable of providing a static map between inputs and outputs (i.e. they have no way of internally representing the dynamics of a nonlinear system). In contrast, RNN may be very difficult (and take a long time) to train stably; however, the recurrent connections of a RNN form a dynamical system. It is this dynamical nature of RNN that allows them to capture the dynamics of a nonlinear system, which makes them more applicable to nonlinear system modelling and control. Examples of RNN are Hopfield networks, Elman networks, liquid state machines, and echo state networks, the last-mentioned being the focus of this paper. For a review of these RNN and the problems associated with training them, see either [1], [2], or [3].

The recent development of echo state networks[4] (ESNs) provides a class of RNN that alleviates the problem of training, but the design methodology of ESNs is still not fully understood. ESNs are characterized by their ability to uniquely map a temporal input history to an “echo state”. An ESN that has this characteristic is said to have the *echo state property*. Currently, the echo state property may be verified from two sufficient conditions, namely one for the existence of echo states for all inputs and one for the non-existence of echo states for certain inputs. This paper reformulates and further develops these conditions into separate necessary and sufficient condition for the existence of echo states for all inputs. As mentioned in [4], the current sufficient condition for the existence of echo states appears, in practice, to be rather restrictive. This problem is addressed by deriving a new sufficient condition that is less conservative. Specifically, a result that is well known in the *Robust Controls* community is used to reduce the conservatism and in some cases make the bounds tight (i.e. provide a single bound that is both necessary and sufficient). This paper concludes with some simple simulations to demonstrate the improvement that can be achieved by using the new sufficient condition.

II. NOTATION

Let $\mathbb{F} = \mathbb{C}$ or $\mathbb{R}$ be either the field of complex or real numbers, respectively. For any square matrix $W \in \mathbb{F}^{n \times n}$, let $\bar{\sigma}(W)$ denote the largest singular value and let $\rho(W) \doteq \max\{|\lambda| : \lambda \text{ is an eigenvalue of } W\}$ denote the spectral radius of W . For any square matrix $D \in \mathbb{F}^{n \times n}$, let $\mathcal{N}(D)$ denote the null space of D . A column vector $x_k \in \mathbb{F}^n$ is a vector of length n at discrete-time instance k (i.e. $x_k \equiv [x_{1k}, x_{2k}, \dots, x_{nk}]^T$, where $[\cdot]^T$ is the standard matrix transpose). The set U^{-N} contains all inputs of N previous time steps $\{u_{k-N}, \dots, u_{k-1}, u_k\}$ that are bounded by some uniform constant; it is therefore compact. Similarly, the sets $U^{-\infty}$ and $U^{+\infty}$ contains all left and right infinite input sequences $\{\dots, u_{k-1}, u_k\}$ and $\{u_k, u_{k+1}, \dots\}$, respectively, which are bound by some uniform constant; they are compact too. Certain scaling matrices will be used. The classes of these matrices will be denoted as $\mathcal{D} \equiv \mathbb{F}^{n \times n}$ and $\bar{\mathcal{D}} = \{\text{diag}(\delta_1, \dots, \delta_n), \delta_i \in \mathbb{F} \text{ for } i = 1, \dots, n\}$. To avoid confusion, D_δ will be used to represent diagonal matrices in $\bar{\mathcal{D}}$ and D will be used to represent full matrices in $\mathcal{D}$. I_n is an $n \times n$ identity matrix. The subscript for this will be dropped when the dimensions are obvious from context. The function $f(\cdot)$ applied to a vector is an element-by-element operation (i.e. $f(x) \equiv [f(x_1), f(x_2) \dots f(x_n)]^T$). As a final note, all of the results are valid for $\mathbb{F} = \mathbb{C}$ or $\mathbb{R}$; however, in the context of an ESN, it is typically assumed that $\mathbb{F} = \mathbb{R}$.

III. PRIOR WORK

An ESN is a multilayered network that consists of a feed-forward input layer, a recurrently connected hidden layer, and a feed-forward output layer. This is shown in Fig. 1.

M. Buehner is a Graduate Student in the Electrical and Computer Engineering Department, Colorado State University (CSU), Fort Collins, Colorado, E-mail: mbuehner@engr.colostate.edu.

P. Young is an Associate Professor in the Electrical and Computer Engineering Department, Colorado State University (CSU), Fort Collins, Colorado, E-mail: pmy@engr.colostate.edu.

Fig. 1. Echo State Network Architecture.

The input layer weights and the hidden layer weights are fixed (i.e. time-invariant) and only the output layer weights are trained (i.e. time-varying). This allows for a neural network that has the dynamic modelling capabilities from the (stable) recurrent connections in the hidden layer and the stable adaptive capabilities from the feed-forward network in the output layer. As will be shown, the echo state property holds for ESNs with asymptotically stable recurrent connections.

The input layer (with weights $W_{in} \in \mathbb{R}^{n \times m}$) is used to map the lower dimensional inputs $u_k \in \mathbb{R}^m$ to a larger dimension $W_{in}u_k \in \mathbb{R}^n$ (i.e. $n > m$). In practice, the weights W_{in} do not appear to have an impact on the performance[4]. The echo states x_k are generated from the input layer and the recurrently and interconnected neurons in the hidden layer. Specifically, for a given an input vector u_k , the transition from one (echo) state vector x_{k-1} to the next (echo) state vector x_k is defined to be

$$x_k = T(x_{k-1}, u_k) = f(W_{in}u_k + Wx_{k-1}), \quad (1)$$

where $W \in \mathbb{R}^{n \times n}$ is a square matrix (containing the hidden layer weights) that is applied to the state vector x_{k-1} , and $f(\cdot)$ is a nonlinear “squashing” function that is applied to every element of the vector $W_{in}u_k + Wx_{k-1}$. The state vectors x_k are fed into the final output layer. In the general statement of an echo state network, feedback connections from the output layer to the hidden layer are allowed; however, in most applications, these feedback connections are not used. The theory for dealing with these connections will not be explored here.

The hidden layer weight matrix (W) is a randomly generated sparse matrix (i.e. usually only 5% to 20% of the entries are non-zero) that is used to generate a (random) basis for the echo states x_k . The properties of this matrix may be used to determine if an ESN has the echo state property. The definition of an echo state network (i.e. an ESN that has the echo state property) is given next.

Definition (Jaeger[4]) Assume standard compactness conditions (i.e. the inputs u_k and states x_k come from the compact sets $U^{-\infty}$ and $X^{-\infty}$, respectively). Assume that the network has no output feedback connections. Then, the ESN has echo states (i.e. the echo state property) if every echo state vector x_k is uniquely determined for every left infinite input sequence $u^{-\infty} \in U^{-\infty}$.

This definition implies that nearby echo states must represent similar input histories. In turn, this means that the echo states should depend more heavily on the most recent inputs and states. Intuitively, this means that the state-space is not disjoint, and the echo states represent the current dynamics, or state, of the system the ESN is modelling. In [4], this property is shown to be satisfied by requiring the echo states to have a certain convergence property. This is stated more formally in the following discussion.

In Jaeger’s paper[4], the existence of echo states may be verified in terms of separate necessary and sufficient conditions on the (square) hidden layer weight matrix $W \in \mathbb{R}^{n \times n}$. The necessary condition is $\rho(W) < 1$ and the sufficient condition is $\bar{\sigma}(W) < 1$. In [4], the necessary condition is stated as a sufficient condition for the nonexistence of echo states when $\rho(W) > 1$. The reason for this constraint is that if the underlying linear system is unstable, then the nonlinear system (resulting from the application of the squashing function) will also exhibit instability. From this point of view, the sufficient condition for the nonexistence of echo states is really for $\rho(W) \geq 1$. The necessary condition for the existence of echo states results from this. From a systems point of view, this requires that the nonlinear recurrent system be locally asymptotically stable at the

origin. For global asymptotic stability, a more restrictive (sufficient) condition is required¹. The original proof of the sufficient condition (taken from [4]) is outlined next.

Let x_k and $\tilde{x}_k$ be two distinct state vectors and $y_k = x_k - \tilde{x}_k$. From an equivalent definition, echo states exist if the states x_k and $\tilde{x}_k$ satisfy the convergence property $\|y_k\| \rightarrow 0$ as $k \rightarrow \infty$ for all right infinite input sequences $u^{+\infty} \in U^{+\infty}$. Note that in this definition, the vector norm is not specified; however, in Jaeger's proof, the standard Euclidean norm (i.e. the 2-norm) is used. Since all finite dimensional norms are equivalent, proving convergence in the 2-norm guarantees convergence in every other (finite dimensional) norm. Jaegers's original sufficient condition is now (re)stated as Theorem 1.

Theorem 1 (Jaeger[4]): Let an echo state network have a fixed internal weight matrix $W \in \mathbb{R}$ and let $f(x) = \tanh(x)$. If $\bar{\sigma}(W) < 1$, then the network has the echo state property, i.e. $\lim_{k \rightarrow \infty} \|y_k\|_2 = 0$ for all right infinite input sequences $u^{+\infty} \in U^{+\infty}$.

Proof:

$$\begin{aligned}
\|y_{k+1}\|_2 &= \|x_{k+1} - \tilde{x}_{k+1}\|_2 \\
&= \|T(x_k, u_k) - T(\tilde{x}_k, u_k)\|_2 \\
&= \|f(W^{\text{in}}u_k + Wx_k) - f(W^{\text{in}}u_k + W\tilde{x}_k)\|_2 \\
&\leq \|(W^{\text{in}}u_k + Wx_k) - (W^{\text{in}}u_k + W\tilde{x}_k)\|_2 \\
&= \|Wx_k - W\tilde{x}_k\|_2 \\
&= \|W(x_k - \tilde{x}_k)\|_2 \\
&\leq \|W\|_2 \|y_k\|_2 \\
&= \bar{\sigma}(W) \|y_k\|_2,
\end{aligned} \tag{2}$$

where W satisfies the contraction property $\bar{\sigma}(W) < 1$. Note that this clearly implies the required convergence property, namely $\lim_{k \rightarrow \infty} \|y_k\|_2 = 0$. ■

Remark Note that eqn (2) assumes that the squashing function will shrink *every* element of the vectors $W^{\text{in}}u_k + Wx_k$ and $W^{\text{in}}u_k + W\tilde{x}_k$ towards zero. Therefore, the difference of the “un-squashed” version will have a larger norm than the difference of the “squashed” version. In the proof given in [4], the squashing function is assumed to be $f(x) = \tanh(x)$; however, it is mentioned that any function satisfying the (element-wise) Lipschitz condition $|f(v) - f(z)| \leq |v - z| \forall v, z \in \mathbb{R}$ will do. As a side note, if $f(\cdot)$ is differentiable, then the Lipschitz condition is equivalent to $|f'(v)| \leq 1 \forall v \in \mathbb{R}$.

This proof yields a rather conservative sufficient condition[4]. In this paper, a less restrictive sufficient condition is derived by considering a different norm.

IV. THE WEIGHTED OPERATOR NORM

In linear algebra, it is a well known fact that there exists an operator norm (sometimes referred to as an induced norm) for a matrix that is arbitrarily close to the spectral radius of the matrix. This is summarized in Lemma 2.

Lemma 2: For every matrix $W \in \mathbb{F}^{n \times n}$ and for every $\epsilon > 0$, there exists an operator norm $\|\cdot\|_D$ such that

$$\rho(W) \leq \|W\|_D \leq \rho(W) + \epsilon \tag{4}$$

Proof: See [5]. ■

Remark This operator norm is achieved by choosing an appropriate weighted operator norm, namely $\|W\|_D = \|DWD^{-1}\|$ with $D \in \mathbb{F}$ nonsingular, that is specific to the matrix W . This weighted operator norm does not depend on the underlying norm used (e.g. any of the p -norms such as $p = 1, 2$, or ∞), but rather on the weighting matrix $D \in \mathbb{F}$ that is selected based on the matrix W . Note that all finite-dimensional norms are equivalent. Therefore, choosing a different norm may require a different D ; however, the property will hold for any chosen norm. For computational reasons, the 2-norm will be used.

A. The Vector D -Norm

The D -norm of a vector $x \in \mathbb{F}^n$ is defined to be $\|x\|_D = \|Dx\|$, where $D \in \mathbb{F}^{n \times n}$ is non-singular and $\|\cdot\|$ is a vector norm (e.g. one of the p -norms). It is easy to show that $\|\cdot\|_D$ is in fact a vector norm provided D is nonsingular [6]. In this paper, the D -norm will be defined in terms of the weighted 2-norm as $\|x\|_D = \|Dx\|_2$, where D is an arbitrary nonsingular matrix to be chosen later.

¹For an ESN, bounded-input bounded-out (BIBO) stability is satisfied trivially, since outputs from the squashing functions are bounded for all inputs. Here, the echo state property is defined in terms of an global asymptotic stability requirement on the echo state vectors.

B. The Matrix Operator D -Norm

Let $x = D^{-1}y$, where $y \in \mathbb{F}^n$. Then, the induced D -norm of a matrix $W \in \mathbb{F}^{n \times n}$ is given as:

$$\begin{aligned}\|W\|_D &= \sup_{x \neq 0} \frac{\|Wx\|_D}{\|x\|_D} \\ &= \sup_{x \neq 0} \frac{\|DWx\|_2}{\|Dx\|_2} \\ &= \sup_{y \neq 0} \frac{\|DWD^{-1}y\|_2}{\|y\|_2} \\ &= \bar{\sigma}(DWD^{-1}),\end{aligned}\tag{5}$$

where $\bar{\sigma}(DWD^{-1})$ is the largest singular value of the matrix DWD^{-1} . Since D is nonsingular, $\mathcal{N}(D) = \mathcal{N}(D^{-1}) = \{0\}$, so $y = 0$ if and only if $x = 0$. Therefore the constraint $y \neq 0$ is equivalent to $x \neq 0$. The last equality in eqn (5) follows from the definition of the induced 2-norm of a matrix.

C. Minimizing the Matrix Operator D -Norm

Since D is arbitrary, it may be chosen such that $\|W\|_D = \bar{\sigma}(DWD^{-1})$ satisfies Lemma 2 for a given ϵ . If D is allowed to have full structure, then

$$\inf_{D \in \mathcal{D}} \bar{\sigma}(DWD^{-1}) = \rho(W),\tag{6}$$

where infimum is used instead of minimum since D (or D^{-1}), in many cases, will be approaching a singular matrix.

If $\mathcal{D}$ is a set of matrices that has some structure imposed upon it, say the set $\overline{\mathcal{D}}$, then $\|W\|_{D_\delta} = \bar{\sigma}(D_\delta W D_\delta^{-1})$ will not necessarily approach the spectral radius of W . Instead, the following relationship holds.

$$\rho(W) \leq \inf_{D_\delta \in \overline{\mathcal{D}}} \bar{\sigma}(D_\delta W D_\delta^{-1}) \leq \bar{\sigma}(W)\tag{7}$$

In eqn (7), the upper bound is obvious since $D_\delta = I$ is always an option. For general W , taking the infimum over all possible $D_\delta \in \overline{\mathcal{D}}$ will result in a measure that is strictly less than the $\bar{\sigma}(W)$ and strictly greater than $\rho(W)$. However, there are classes of matrices for which the lower bound of eqn (7) is exact. This leads to the following theorem.

Theorem 3: Let $W \in \mathbb{F}^{n \times n}$ be in one of the following two classes:

- 1) normal matrices, and
- 2) upper and lower triangular matrices.

Then, there exists a $D_\delta \in \overline{\mathcal{D}}$ such that $\|W\|_{D_\delta} = \rho(W) + \epsilon$ for all $\epsilon > 0$.

Proof:

Class 1: The first class is the easiest to prove since the singular values of a normal matrix are equal to the absolute values of its eigenvalues. Therefore, the maximum singular value and spectral radius are equal.

Class 2: In contrast, the gap between the spectral radius and maximum singular value of a triangular matrix may be arbitrarily large, but the operator D_δ -norm can always be made arbitrarily close to the spectral radius. To see this, let $\delta > 0$, $W \in \mathbb{F}^{n \times n}$ be upper triangular, $D_\delta = \text{diag}(1, \delta, \delta^2, \dots, \delta^{n-1})$, and $D_\delta^{-1} = \text{diag}(1, \frac{1}{\delta}, \frac{1}{\delta^2}, \dots, \frac{1}{\delta^{n-1}})$. Then,

$$D_\delta W D_\delta^{-1} = \begin{pmatrix} w_{1,1} & \frac{1}{\delta} w_{1,2} & \cdots & \frac{1}{\delta^{n-1}} w_{1,n} \\ 0 & w_{2,2} & \ddots & \vdots \\ \vdots & \ddots & \ddots & \frac{1}{\delta} w_{n-1,n} \\ 0 & \cdots & 0 & w_{n,n} \end{pmatrix}\tag{8}$$

Using a limiting argument yields

$$\lim_{\delta \rightarrow \infty} \bar{\sigma}(D_\delta W D_\delta^{-1}) = \bar{\sigma} \begin{pmatrix} w_{1,1} & 0 & \cdots & 0 \\ 0 & w_{2,2} & \ddots & \vdots \\ \vdots & \ddots & \ddots & 0 \\ 0 & \cdots & 0 & w_{n,n} \end{pmatrix} = \max_{1 \leq k \leq n} |w_{k,k}| = \rho(W),\tag{9}$$

which from Lemma 2 is the smallest that any operator norm may approach. ■

Remark For a lower triangular matrix, the same argument would be used with the limit $\delta \rightarrow 0$.

Remark In this derivation, D^{-1} is approaching a (rank one) singular matrix, but a finite D may be chosen that is arbitrarily close to the infimum as stated in Lemma 2.

Remark Theorem 3 also holds for matrices that may be permuted to triangular matrices by swapping the matching rows and columns (e.g. swapping rows 3 and 5 followed by swapping columns 3 and 5). In the context of an ESN, this amounts to a re-labelling of the recurrently connected neurons.

The results from this section will be used to obtain a tighter (and in some cases exact) bound for the echo state property.

V. A NEW SUFFICIENT CONDITION FOR THE ECHO STATE PROPERTY

In this section, a new sufficient condition is derived using the results from the previous section.

Theorem 4: Let an echo state network have a fixed internal weight matrix $W \in \mathbb{R}$ and assume the squashing function $f(\cdot)$ satisfies the (element-wise) Lipschitz condition $|f(v) - f(z)| \leq |v - z| \forall v, z \in \mathbb{R}$. If $\inf_{D_\delta \in \overline{\mathcal{D}}} \bar{\sigma}(D_\delta W D_\delta^{-1}) < 1$, then the network has the echo state property, i.e. $\lim_{k \rightarrow \infty} \|y_k\|_{D_\delta} = 0$ for all right infinite input sequences $u^{+\infty} \in U^{+\infty}$.

Proof:

$$\begin{aligned} \|y_{k+1}\|_{D_\delta} &= \|x_{k+1} - \tilde{x}_{k+1}\|_{D_\delta} \\ &= \|T(x_k, u_k) - T(\tilde{x}_k, u_k)\|_{D_\delta} \\ &= \|f(W^{\text{in}}u_k + Wx_k) - f(W^{\text{in}}u_k + W\tilde{x}_k)\|_{D_\delta} \\ &\leq \|(W^{\text{in}}u_k + Wx_k) - (W^{\text{in}}u_k + W\tilde{x}_k)\|_{D_\delta} \end{aligned} \quad (10)$$

$$\begin{aligned} &= \|Wx_k - W\tilde{x}_k\|_{D_\delta} \\ &= \|W(x_k - \tilde{x}_k)\|_{D_\delta} \\ &\leq \|W\|_{D_\delta} \|y_k\|_{D_\delta} \\ &= \bar{\sigma}(D_\delta W D_\delta^{-1}) \|y_k\|_{D_\delta}. \end{aligned} \quad (11)$$

where $\inf_{D_\delta \in \overline{\mathcal{D}}} \bar{\sigma}(D_\delta W D_\delta^{-1}) < 1$ satisfies the contracting property. Note that this clearly implies the required convergence property, namely $\lim_{k \rightarrow \infty} \|y_k\|_{D_\delta} = 0$. ■

Remark Equation (10) is the reason for using the set of diagonal scaling matrices $\overline{\mathcal{D}}$, which follows from removing the squashing function. If D_δ was allowed to have full structure (i.e. if $D \in \mathcal{D}$ was used), this inequality will not hold for all state vectors x_k and $\tilde{x}_k$. However, the diagonally structured D_δ ensures that every element of the “squashed” version will be less than the “un-squashed” version and hence the D_δ -norm will be less. Therefore, a sufficient condition for the existence of echo states is $\inf_{D_\delta \in \overline{\mathcal{D}}} \bar{\sigma}(D_\delta W D_\delta^{-1}) < 1$.

Remark Since all finite dimensional norms are equivalent, $\lim_{k \rightarrow \infty} \|y_k\|_{D_\delta} = 0$ implies that $\lim_{k \rightarrow \infty} \|y_k\|_2 = 0$. Therefore, the original echo state property is satisfied by the new constraint.

Corollary 5: If W is a normal matrix or a (permuted) triangular matrix, then $\rho(W) < 1$ is both a necessary and sufficient condition for the existence of echo states for all inputs u_k .

Proof: The proof of this follows from Theorem 3. ■

Remark If $W \in \mathbb{F}$ is triangular, then $\rho(W) < 1$ is identical to $\max |\text{diag}(W)| < 1$.

Remark The stability requirement $\inf_{D_\delta \in \overline{\mathcal{D}}} \bar{\sigma}(D_\delta W D_\delta^{-1}) < 1$ is equivalent to a strong form of Lyapunov stability[7].

Since D must be diagonal, this new bound is not tight in the sense that it is not necessarily equivalent to $\rho(W) < 1$ for any arbitrary W , but it is considerably less conservative than the bound $\bar{\sigma}(W) < 1$.

The reason for using the operator 2-norm (as the underlying norm) is that there exists commercial software for minimizing $\|DW D^{-1}\|_2$ when D must be a structured matrix. In fact, this is a well studied problem in the robust control community where $\|DW D^{-1}\|_2$ is an upper bound for the structured singular value of W (usually denoted as $\mu(W)$). For more information on robust controls and the structured singular value, see [8]. Using (for example) MATLAB’s *μ -Analysis and Synthesis Toolbox*[9], the infimum of $\|W\|_{D_\delta}$ may be calculated using the command

`[muUB, muLB] = mu(W);`

Here, $\text{muUB} = \inf_{D_\delta \in \overline{\mathcal{D}}} \bar{\sigma}(D_\delta W D_\delta^{-1})$. Therefore, if $\text{muUB} < 1$, the ESN has the echo state property. Note that if $\text{muUB} \geq 1$, a new ESN may be defined with internal matrix $\tilde{W} = \frac{.99}{\text{muUB}} W$, which will satisfy the new sufficient condition, and hence, have the echo state property.

TABLE I
RESULTS OF ESN SIMULATIONS FOR UNSTRUCTURED W .

	$\bar{\sigma}(W)$	$\ W\ _{D_\delta}$	$\rho(W)$	mse
$\bar{\sigma}(W) < 1$	0.99	0.7390	0.5303	0.0136
$\ W\ _{D_\delta} < 1$	1.3099	0.99	0.7056	0.0127
$\rho(W) < 1$	1.9111	1.4913	0.99	0.0148

VI. TRIANGULAR VS. UNSTRUCTURED WEIGHT MATRICES

In order to demonstrate the usefulness of the new sufficient condition, an ESN was used to model a mass-spring-damper system. The details of this experiment may be found in [10]. The basic ESN architecture was two inputs, namely position and velocity, and one output, namely acceleration. The hidden layer contained 20 neurons (i.e. $W_{in} \in \mathbb{R}^{20 \times 20}$), and the output layer was trained to estimate the acceleration based on the echo states generated from the previous position and velocity inputs. After the output layer was trained, a test set (which was different than the training set) of (position and velocity) inputs were used to assess the various ESNs. Performance for each ESN was assessed in terms of the mean-squared error (mse) between the predicted acceleration and the theoretical acceleration. The next two sections summarize the results from simulating ESNs with full and triangular hidden layer weight matrices.

A. ESN with Unstructured Hidden Weight Matrices

For these experiments, the hidden layer weight matrix W had a density of 20%, which meant that it had $0.2 \times 20^2 = 80$ non-zero entries. The ESNs were simulated (with W scaled tight against the bound) using the old sufficient condition, new sufficient condition, and the necessary condition only. Recall that satisfying the sufficient conditions guarantees satisfying the necessary condition but not vice versa. In these experiments, W was created from sampling 80 entries of a precalculated random matrix $W_{rand} \in \mathbb{R}^{20 \times 20}$, which was fixed for all experiments. The resulting W matrices were then scaled to be tight against the bound being used. In each of these experiments, ten random samplings were used and their results averaged. The averaged results are summarized in Table I.

For this particular experiment, the ESNs that used the new sufficient condition worked the best; however, the performance of a given ESN depends strongly on the topology used. The topic of topology is beyond the scope of this paper; however, a more detailed explanation of the dependence of an ESN on topology may be found in [10]. While the bound $\rho(W) < 1$ is generally used in practice[4], this scaling is “dangerous” in the sense that it does not guarantee the echo state (global stability) property for all allowable inputs. The interest here stems from the fact that a larger scaling parameter may be used and still guarantee global stability of an ESN. In most cases tested, the general topology (usually non-optimal) performed the best when the hidden layer weight matrix was scaled to be tight against the new sufficient condition bound.

B. ESNs with Structured Hidden Weight Matrices

The same experiment (as mentioned above) was used to test the usefulness of both normal and triangular hidden layer weight matrices W . In these experiments, the normal matrix was a symmetric matrix with 80 non-zero entries, and the triangular matrix was an upper triangular matrix with 40 non-zero entries. Since the rank of W appears to have an impact on the performance of an ESN[10], the upper triangular W was forced to have full rank by assigning non-zero values along its main diagonal. The results of these experiments are presented in Table II.

TABLE II
RESULTS OF ESN SIMULATIONS FOR STRUCTURED W .

	$\bar{\sigma}(W)$	$\ W\ _{D_\delta}$	$\rho(W)$	mse
symmetric	0.99	0.99	0.99	0.0143
triangular	1.6698	0.99	0.99	0.0155

These structured weight matrices produced similar results to the unstructured matrices that satisfied the necessary condition only. However, these matrices, when scaled so that $\rho(W) < 1$, are guaranteed to provide stable ESNs where the unstructured matrices are not. For a more detailed study of how to design and train echo state networks, see [10] and [4].

VII. CONCLUDING REMARKS

The echo state property is directly related to the asymptotic stability of the dynamical system formed in the hidden layer of an ESN. Even though ESNs appear to perform well in practice with only local stability (i.e. only the necessary condition is satisfied) [3], [4], [11], [10], [12], there is no guarantee that they will always perform well on a physical system. In order to make a more concrete statement about the performance of an ESN for all inputs (i.e. to guarantee that echo states exist for all inputs), a more rigorous global asymptotic stability requirement is needed. While a (rather restrictive) requirement existed

for guaranteeing global asymptotic stability, this paper presented a new bound that is less restrictive than the original bound and for some topologies provides exact analysis. Note that this paper does not address the design methodology of an ESN (for that see [10]). Instead, it provides a test for global asymptotic stability (and hence the existence of echo states for all inputs) which utilizes linear algebra optimization problems with readily available commercial software.

VIII. ACKNOWLEDGEMENTS

The authors would like to thank Keith Bush for his helpful discussions and for providing the code used to test the echo state networks. The authors would also like to thank the National Science Foundation for providing funding for this project under award 0245291.

REFERENCES

- [1] J. J. Steil, *Input-Output Stability of Recurrent Neural Networks*. Göttingen: Cuvillier Verlag, 1999, (Also: Phd.-Dissertation, Faculty of Technology, Bielefeld University, 1999). [Online]. Available: <http://www.techfak.uni-bielefeld.de/~jsteil/publications.html>
- [2] A. F. Atiya and A. G. Parlos, "New results on recurrent network training: Unifying the algorithms and accelerating convergence," *IEEE-NN*, vol. 11, no. 3, p. 697, May 2000. [Online]. Available: citeseer.ist.psu.edu/atiya00new.html
- [3] P.-G. Plöger, A. Arghir, T. Günther, and R. Hosseiny, "Echo state networks for mobile robot modeling and control," in *RoboCup*, 2003, pp. 157–168. [Online]. Available: <http://springerlink.metapress.com/openurl.asp?genre=article&issn=0302-9743&volume=3020&page=157>
- [4] H. Jaeger, "The echo state approach to analyzing and training recurrent neural networks," German National Research Center for Information Technology, Tech. Rep. 148, 2001.
- [5] R. Horn and C. R. Johnson, *Matrix Analysis*. Cambridge University Press, 1985.
- [6] L. N. Trefethen and D. Bau III, *Numerical Linear Algebra*. SIAM, Philadelphia, 1997.
- [7] A. Packard and J. C. Doyle, "The complex structured singular value," *Automatica*, vol. 29, no. 1, pp. 71–109, 1993.
- [8] S. Skogestad and I. Postlethwaite, *Multivariable Feedback Control*. John Wiley and Sons Ltd., West Sussex, England, 1996.
- [9] G. J. Balas, J. C. Doyle, K. Glover, A. Packard, and R. Smith, *μ -Analysis and Synthesis TOOLBOX User's Guide*. The Math Works, Inc., Natick, Mass., 1996.
- [10] K. Bush and B. Tsendsjav, "Improving the richness of echo state features using next ascent local search," in *Proceedings of the Artificial Neural Networks In Engineering Conference*, St. Louis, MO, November 2005.
- [11] K. Bush and C. Anderson, "Modeling reward functions for incomplete state representations via echo state networks," in *Proceedings of the International Joint Conference on Neural Networks*, Montreal, Quebec, July 2005.
- [12] K. Ishii, T. Van Der Zant, V. Becanovic, and P. G. Plöger, "Identification of motion with echo state network," in *Oceans'04 (MTS/IEEE)*, Kobe, Japan, Nov 2004.
- [13] J. W. Demmel, *Applied Numerical Linear Algebra*. SIAM, Philadelphia, 1997.