

Assessment of Damage in an 8-Oscillator Circuit Using Dynamical Forcing

Linda Moniz,* Thomas Carroll, and Louis Pecora

Code 6340 Naval Research Laboratory, Washington,

DC 20375 USA. Phone (202) 767 6259, fax (202) 767 1697

Michael Todd

Department of Structural Engineering, University of California,

San Diego, La Jolla, CA 92093-0085 USA

(Dated: May 20, 2003)

Abstract

We employ chaotic interrogation of a circuit simulation of a structure in order to test for damage to the structure. The circuit simulation provides a realistic test of our attractor-based method and permits close control over parameters in the structure. In this circuit, simulating an 8-degree of freedom spring-mass system, we were able to detect changes of as little as 2% in the coupling between two oscillators in the circuit. This corresponded to detection of a 2% loss in stiffness to one spring in the modelled system.

PACS numbers: 02.40-k,02.50-r,05.45Tp,05.45Xt

Keywords: damage detection, attractor, lyapunov exponents, continuity test, filtered Lorenz signal, multiple time-series embedding

*Electronic address: moniz@anvil.nrl.navy.mil; This project was supported by an ONR-ASEE Fellowship and a 6.1 Navy ARI.

I. INTRODUCTION

Simulation of a physical system with an electrical circuit affords the experimentalist control over the parameters of the system that are nearly as good as in a numerical simulation and at the same time provides insight into the behavior of data analysis algorithms with real experimental data. Typically, output from a circuit simulation is also low noise. This allows the analyst to concentrate on the analysis of the data and method without the additional problem of dealing with noise.

Here we develop an 8-oscillator circuit simulation of a spring-mass damper system. We simulate incremental changes in stiffness to one spring near the fixed end of the system and force the circuit at the opposite end with a Lorenz signal. The system acts as a linear filter of the signal; this linear filter changes with changes in stiffness. It has been shown both theoretically (see, e.g. [5], [12]) and experimentally ([1, 9]) that in some cases (e.g. IIR filters) the filtering of a chaotic signal changes the dimension of the attractor that is reconstructed from the signal. Todd, Nichols et al showed that it is possible to extract various features from the reconstructed attractors from chaotically forced structures in order to detect damage ([7, 14]). Here, we test the limits of detection by attempting to detect very low levels of damage in only one location in a structure. We introduce a new statistical test, based on the continuity test of Pecora, Carroll and Heagy [10] to detect geometric changes in the reconstructed attractors and test to see how this statistic scales with the level of damage.

In our experiment we record the output at each oscillator and use multivariate time-series analysis to embed the time series from various change scenarios. We note that although previous multivariate embedding methods were tailored to time series that were weakly coupled ([4]), here we use very strongly coupled time series, a result of the forcing of the structure. We develop strategies for finding the dimension of the embedding using a series of false-nearest-neighbor tests.

Our statistical test determines if there is evidence that a continuous map exists between the reconstructed attractor from a baseline system to that of each change scenario. Using this test, we are able to detect a 2% change in stiffness from the baseline scenario. We also observe incremental changes in the statistic that parallel the incremental changes in the circuit.

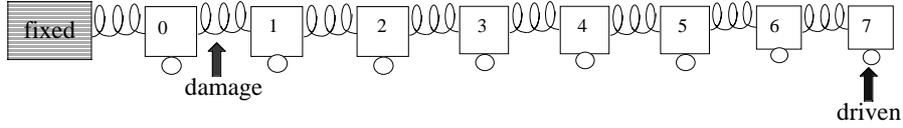


FIG. 1: 8 degree-of-freedom Spring Mass System

II. EXPERIMENTAL DESIGN

In order to test damage detection algorithms in a controlled and well characterized system, we built a simple analog electronic circuit to simulate a set of masses connected by springs (see figure 1). The circuit consisted of 8 underdamped oscillators coupled in a line. We based the oscillator circuits on the following model:

$$\begin{aligned} \frac{dx_i}{dt} &= \tau_i y_i \\ \frac{dy_i}{dt} &= -\tau_y [\gamma(2y_i - y_{i+1} - y_{i-1}) + \alpha(2x_i - x_{i-1} - x_{i+1})] \end{aligned} \quad (1)$$

The average time constant τ_i for the 8 oscillators is 10^4 s, but the individual τ_i 's vary by 20% from this average to simulate variations in the object being tested. The other time constant, τ_y , is constant at 10^4 . Damage is simulated by changing the stiffness constant α from its normal value of 1, while the damping constant γ is set to .05.

The circuit for one oscillator is shown in figure 2. The value of x_i is measured at the location marked X , and y_i is measured at Y . The oscillator time constants are $\tau_i = 1/R_{\tau_i}C$, where $C = .001\mu F$ and R_{τ_i} is given in table 3. We simulate damage by increasing the values of $R_{x(i+1)}$ and $R_{x(i-1)}$ from their normal values of $10k\Omega$. The signals from the other oscillators are input at the coupling points indicated.

We are simulating an array of masses and springs where one end is fixed and the other end is driven. The driven oscillator is oscillator 7, while the oscillator at the fixed end is oscillator 0. To simulate the fixed end, we set x_{-1} and y_{-1} equal to 0. At the driven end, $x_8 = 0$ and y_8 (the driving signal) is a computer generated Lorenz signal.

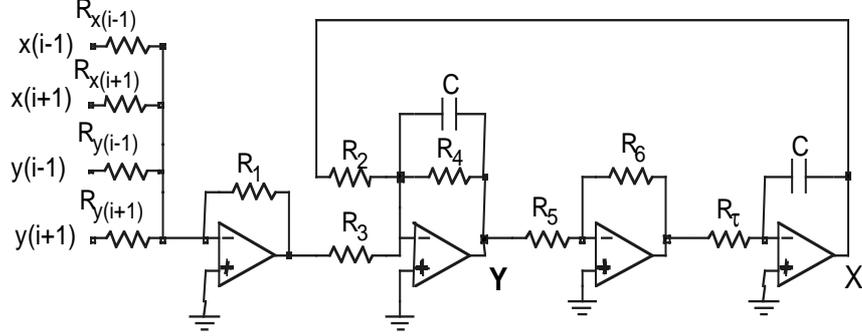


FIG. 2: Single underdamped oscillator circuit used in the experiment. $R_{x(i+1)}$ and $R_{x(i-1)}$ have nominal values of $10k\Omega$, but their values can be increased to simulate damage. $R_{y(i\pm 1)} = 200k\Omega$, $R_1 = 10k\Omega$, $R_2 = 50k\Omega$, $R_3 = 100k\Omega$, $R_4 = 1M\Omega$, $R_5 = 100k\Omega$, $R_6 = 100k\Omega$, $C = 0.001\mu F$. R_7 is given in Table 3. All resistance values are $\pm 1\%$.

The Lorenz system used to generate the driving signal is

$$\begin{aligned}
 \frac{dx}{dt} &= 16(y - x) \\
 \frac{dy}{dt} &= 45.92x - xz - y \\
 \frac{dz}{dt} &= xy - 4z
 \end{aligned} \tag{2}$$

These Lorenz equations are numerically integrated with a 4th order Runge-Kutta routine at a time step of .002. The Lyapunov exponents for this Lorenz system are $1.5s^{-1}$, 0 and $-22s^{-1}$. The Lorenz x signal is played back through a digital to analog converter at 110,000 points/s to form the driving signal. At this playback rate, the Lyapunov exponents for the Lorenz signal are multiplied by $(110,000 \text{ points/s})/(500 \text{ points/s})$ so the new set of Lyapunov exponents are $330s^{-1}$, 0 and $-4840s^{-1}$. These numbers may be compared to the real parts of the first 3 eigenvalues for the undamaged circuit (measured experimentally) of $-123s^{-1}$, $-149s^{-1}$ and $-250s^{-1}$.

The signals from the driven oscillator array were digitized at 22,000 points/s.

i	0	1	2	3	4	5	6	7
$R_{\tau_i}(\text{K}\Omega)$	121	100	80.6	90.9	121	110	110	90.9

FIG. 3: Values of Resistor R_{τ_i} for different oscillators.

III. THE SKEW-PRODUCT SYSTEM AND STRUCTURAL RESPONSE TO CHAOTIC EXCITATION

The driven system we describe in section II is a particular skew-product that arises from the filtering of a dynamical system. Here we offer a characterization of these systems which will lead to the geometric test for damage in the circuit. Much of this background information is covered in Davies and Campbell [5], and we follow their notation with minor changes for clarity.

Consider a nonlinear function $f : X \rightarrow X$, where we are interested in the evolution $f^n(x)$ of points $x \in X$, for $n \in \mathcal{Z}$. After a period of transient behavior, the iterates of x under f settle into the attractor of f in phase space.

Now consider applying a linear filter to a nonlinear function. Let the linear filter be the matrix B . Let $X \in \mathfrak{R}^m$ be compact and let the nonlinear function be $f : X \rightarrow X$ be written as $x_n := f^n(x)$. Thus, we have the following slaved dynamical system:

$$\begin{aligned} y_{n+1} &= By_n + x_n \\ x_{n+1} &= f(x_n) \end{aligned} \tag{3}$$

The input function is unchanged by the filter. However, the slaved system $\{y_n\}$, is determined by the filtering of the base signal $\{x_n\}$ (the drive signal). This is an example of an Infinite Impulse Response (IIR) filter; the series y_n is affected by the entire time history of y . In practice, the dynamical system f is not directly observable; we instead have an observation function $H : X \rightarrow \mathfrak{R}^p$. The function H can be thought of as an appropriate coupling function.

The stability of any dynamical system is described by its Lyapunov exponents. In the case of a linear function, the Lyapunov exponents are the logarithms of the moduli of the characteristic values (i.e. the characteristic exponents). We say that such a system is *stable* if all of the characteristic values lie inside the unit circle. In the case of an IIR filter, as

long as B is stable, the Lyapunov exponents of the IIR filter are simply the characteristic exponents of B together with the Lyapunov exponents of f (see, e.g. Davies and Campbell [5]). In a practical situation, all filters provided by linear structures are stable.

Now we can define the skew product system F on the entire space $X \times \mathbb{R}^p$ by the following:

$$F(x, y) = (f(x), By + H(f(x))) \quad (4)$$

We have fulfilled the criteria for lemma 1 of Davies and Campbell [5]. Thus, there exists a continuous function ϕ such that the graph of ϕ represents an unique, globally attracting F -invariant manifold. The function ϕ is a consequence of the skew product and of the stability of the filter B . In our case, the attractor given by the Lorenz function is not a manifold. However, we may still define the function ϕ in the same way. Although we cannot expect that $(x, \phi(x))$ will be a manifold, we can still consider it as an F -invariant geometric object.

Now suppose that we filter the chaotic function f through two different structures, each yielding a stable linear filter. Given these two different stable linear filters, B_1 and B_2 of the *same* chaotic input signal, we may construct functions ϕ_i in the same manner as ϕ , above. We then have two different geometric objects, the graphs of ϕ_1 and ϕ_2 which represent the two skew product systems, call them Y and Z , respectively.

It is natural to define a map between Y and Z using the drive space X and the representation of each point in Y as $(x, \phi_1(x))$ and each point in Z as $(x, \phi_2(x))$. Thus, we present the following. Let $\Psi : Y \rightarrow Z$ be defined by

$$\Psi(x, \phi_1(x)) = (x, \phi_2(x)) \quad (5)$$

This function is well-defined because of the projections on Y and Z which take $(x, \phi_i(x))$ to $x \in X$; because the filter is linear, the functions ϕ_i are one-to-one. We are interested in properties of the function Ψ . In particular, this function shows the relationship between the two geometric objects given by the graphs of the functions ϕ_i .

We would like to develop tests which indicate differences between (time-series reconstructions of) Y and Z which arise because of differences in the filters B_1 and B_2 . Changes in the linear filtering of the input signal will indicate changes in the structure itself.

In an experimental setting, a time series of measurements of some function of the filtered signal is the only available quantity. In this case, the attractor can be reconstructed using a

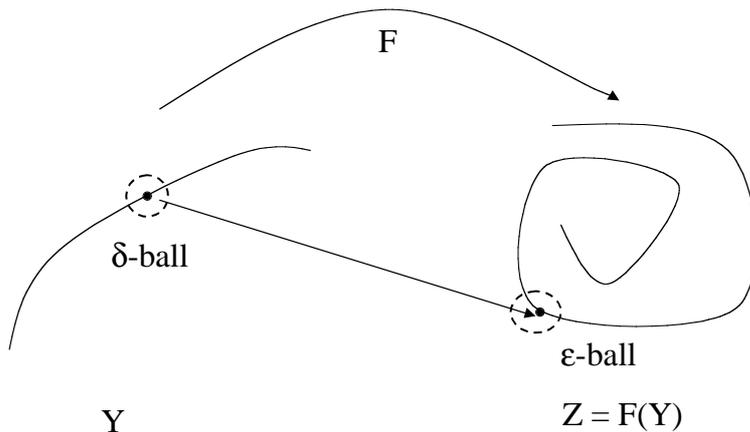


FIG. 4: Points from the δ -ball on the left map to the ϵ -ball on the right

time-series of the observed quantity i.e. a time-delay embedding. It has been shown that the reconstructed attractor is a faithful representation of the original attractor [11]. Because we force our structures with an identical signal, the function Ψ can be constructed implicitly by mapping a time-delay coordinate from the reconstruction of Y to the corresponding (by time) time-delay coordinate in the reconstruction of Z .

We now turn to description of the particular test we use for the changes in Ψ .

IV. THE TEST FOR STATISTICAL CONTINUITY

Given time-delay embeddings of two different geometric objects reconstructed from time series data, it is often important to find a functional relationship between the two objects. For instance, in the presence of noisy data from the reconstruction of one object, can we say if it is essentially identical to another object? Proving or disproving the existence of a continuous function between two such objects can be a powerful tool for analysis of nonlinear behavior. Given a proposed function $F : X \rightarrow Y$ the mathematical definition of continuity at a point $x(t) \in X$ is stated as follows: For all $\epsilon > 0$, $\exists \delta > 0$ such that if $\|x(t_i) - x(t_j)\| < \delta$, then $\|F(x(t_i)) - F(x(t_j))\| < \epsilon$. The geometric meaning of this statement is illustrated in figure 4.

If we let the input to the function be designated as the source and the output designated

as the target, then for an arbitrarily small set in the target, a set in the source can be found for which all points map to the set in the target. Thus, we see that points that are close to each other in the source map to points that are close to each other in the target. The 'closeness' is the relationship between δ and ϵ .

For a theoretical geometric object the analytic definition is clear. However, translating the mathematical $\epsilon - \delta$ definition of continuity to a time-series reconstruction setting raises two important questions:

1. How can potentially noisy, finite data yield a reasonable definition of continuity either at a point or on an entire geometric object?
2. How can such a definition be translated to a meaningful and reliable statistic regarding the absence or presence of a continuous function?

Three clear problems appear when considering continuity in the context of finite data. The first is that ϵ cannot be made to go to zero. Thus, some finite but small ϵ that still indicates continuity will have to be determined. Moreover, for some x , there may be ϵ for which δ can be found even if there is no continuous F . Secondly, only a finite number of points $x \in X$ can be checked for continuity. Lastly, in the presence of noise, even for an obviously continuous F (e.g. an identity function), all points from a δ -ball may not map to the corresponding ϵ -ball. For example, see figure 5. These issues cannot be ignored, but we can create a statistical criterion for continuity that is consistent with the $\epsilon - \delta$ definition.

We begin with two time series reconstructions, denoted Y and Z . The space Y will be denoted the source, Z the target (or image). We formulate the continuity test as follows.

Theoretically, we would choose an ϵ as in the formal definition. We find the ϵ -ball around the point z in the space Z . We then take progressively smaller δ -balls around the corresponding point $y \in Y$ that maps to $z \in Z$ until all of the points in the δ -ball are mapped into the ϵ -ball. However, because of the above issues, we need instead to apply a statistical criterion which will reject or accept the δ -ball as passing the continuity test for this ϵ . For this, we formulate a null hypothesis.

The null hypothesis assumes that for any point $y(t_j)$ in a δ -ball the corresponding point $z(t_j)$ has a probability .5 of being in the ϵ -ball, regardless of the size of the ϵ -ball. If n points are in any δ -ball, the probability of m or more of these points' images in the ϵ -ball must be $< .05$ to reject the null hypothesis.

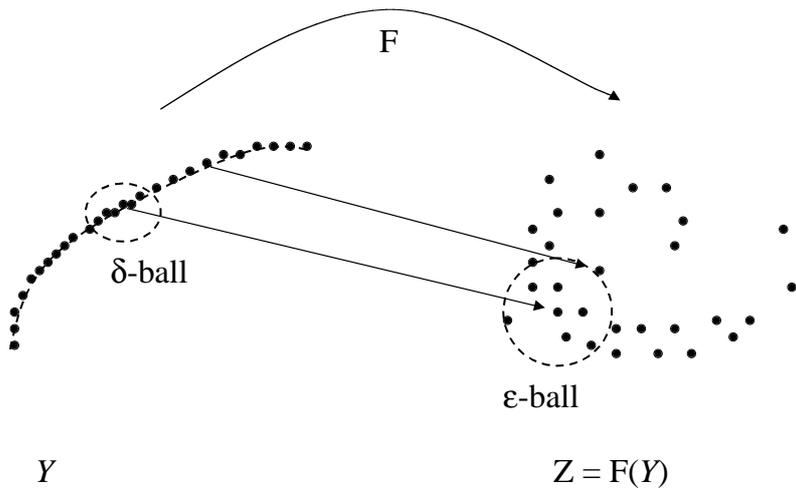


FIG. 5: An ϵ and δ may be found but ϵ may be large. Noise may force some points from the δ -ball to be outside the ϵ -ball.

The null hypothesis essentially assumes that points from the given δ -ball map to points in the ϵ -ball by a coin flip. In order to reject the null (equivalently, to *accept* the δ -ball as passing the continuity test for this ϵ), the probability must lie in the tail of the binomial distribution. Thus, we must have 95% confidence that the points from the δ -ball did *not* map to the ϵ ball by chance.

This differs from the null hypothesis described in [10]. To account for noise, our null hypothesis allows some points from the δ -ball to map outside the ϵ -ball. However, we require that enough points m from the δ -ball map into the ϵ -ball to insure that the probability of m or more points landing in the ϵ -ball by chance (noise) is low. Hence, the possibility that noise can produce evidence of continuity is negligible.

We formulate the statistic to be based not on the acceptance or rejection of the null hypothesis, but on the *minimum* ϵ that can be used to reject the null hypothesis at each point. We call this value ϵ^* .

To compute the continuity statistic, N test points $y(t_i)$ are chosen at random times from Y . This serves to also distribute the points randomly in space. In our implementation, the data are normalized so that the standard deviation of the attractor is $\sigma = 1$. For each test point, initially $\epsilon = \delta = 3\sigma$. The number of points in the δ -ball around the representative point $y(t_i)$ is n . Image points in the ball centered around the point $z(t_i)$ are counted; this

number is m . Then the binomial distribution with parameters $(n, .5)$ is computed to find the cumulative probability of finding m or more image points in the ϵ -ball. If this probability is $< .05$, the null hypothesis is rejected for this point and ϵ is recorded as ϵ^* . Then ϵ is reduced with the same δ . If the null hypothesis is not rejected, δ is reduced. To maintain the 95% confidence interval, there must be at least 5 temporally non-correlated points in the δ -ball. If no ϵ can be found with *any* acceptable δ , we increase the initially allowed ϵ until ϵ^* can be found for all points. Note that ϵ^* for each point represents the *smallest* ϵ for which the null hypothesis is rejected. The average and distribution of ϵ^* is recorded, along with the maximum δ for each ϵ^* .

To detect differences in geometric structures using the continuity test, we compare both the average ϵ^* and the distribution of ϵ^* for the set of representative points. For comparison, we compute the continuity statistics for a known functional relationship in order to see the smallest possible ϵ ; call this value ϵ_0 . If $\epsilon^* \gg \epsilon_0$ for a particular test, it is clear that any functional relationship between the source and target attractors is in question. If ϵ^* is close to ϵ_0 , we examine the distribution for the tests in question and for the known relationship to detect either degradation of a functional relationship or evidence that a continuous functional relationship persists.

For example, if we are testing for continuity of the function Ψ described in section III, we first find ϵ^* values for a function between the attractors reconstructed from two sets of output from an identical *undamaged* circuit. We call this value ϵ_0 . We may define the function Ψ between reconstructed attractors based on the input signal x as in the definition 5. The prerecorded input signal allows us to treat all sets of output data as if they were recorded simultaneously from the same input signal.

We note that ultimately the test on our experimental data will be performed not on the actual geometric object, but on an attractor reconstructed from time series. Thus although theoretically the function Ψ from equation 5 is continuous because it is the composition of continuous functions, we observe that loss of differentiability in a discrete function mimics the loss of continuity (see fig 6). In this case, it may not be possible to tell if the continuity test indicates that the *continuity* or the *differentiability* of the functional relationship changes with damage. In either case, changes in ϵ^* point to changes in the filtering of the chaotic signal.

We emphasize that the continuity statistic is a one-sided statistic. Thus, evidence of a

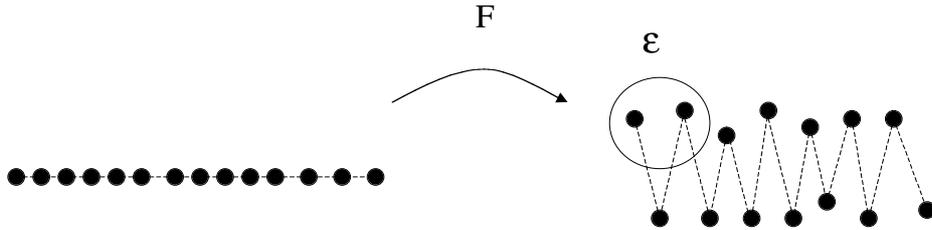


FIG. 6: Loss of differentiability may mimic loss of continuity. For the given ϵ -ball, there is no δ -ball for which all points map to the ϵ -ball.

continuous functional relationship $\Psi_1 : Y \rightarrow Z$ does not imply existence of a continuous function $\Psi_2 : Z \rightarrow Y$. In practice, we compute the continuity statistic using source Y and target Z and then compute the statistics using source Z and target Y . These statistics are considered separately, but we note that in our tests, the continuity tests for both functional directions gave similar statistics.

V. DATA ANALYSIS

Data from these experiments were a set of 8 time series of length 200,000 points. We front-truncated the data by 10,000 points to allow for chaotic transients, and then used the next 80,000 points (a data limitation because of computer memory and speed constraints).

We embedded the 8 time series in a 16-dimensional space, using 2 time-delays per time series. The embedding dimension was determined by using a false-nearest neighbor analysis [8] on all 8 time series. The false-nearest neighbor analysis has been used for determination of embedding dimension for multivariate data by Boccaletti et al[4]. However, in that case the time series were weakly coupled. In our case the time series were strongly coupled via the drive signal. The scheme in [4] was adaptive, in that dimensions were only increased when false near neighbors were found. We adopted a non-adaptive scheme, testing all possibilities for combinations of time series. Because our time series were strongly coupled, it was conjectured that the adaptive scheme might be susceptible to coupling effects between time series. This could cause the adaptive false-near neighbors routine to omit a time series.

The result of this could be an incomplete investigation of embedding dimension. Thus, we investigate embeddings starting with dimension 1 and ending with dimension 24.

We have 8 time series from which we construct d -dimensional delay vectors. For a 1-dimensional embedding, we construct 1-dimensional vectors from a single time series $s_j(t)$. For 2-dimensional vectors we use another time series $s_k(t)$ for the second component of the vector. To further increase the delay to dimensions d with $1 \leq d \leq 8$, we add additional time series. An 8-dimensional time-delay vector looks like (s_1, \dots, s_8) .

For dimensions $d > 8$, we add delays in constructing the vector. One example is $(s_1(t), \dots, s_3(t), s_3(t + \tau), \dots, s_8(t))$. To create embedded vectors of dimension $d > 9$, we continue to add delayed time series $s_j(t + \tau)$ (and in dimensions > 16 , $s_j(t + 2\tau)$) as new components for the vector.

We used the autocorrelation function on all 8 time series to find an appropriate time delay. We used a time-delay window of 30 time steps. This corresponded to an approximately $2/3$ loss of autocorrelation for the time series output by the oscillators at the driven end of the undamaged circuit. Because the autocorrelation delays for the output of the oscillators at the fixed end of the system were much longer (≈ 150 time steps), we used the shorter delay times of the driven end. In a multivariate time series embedding with all output values synchronized with the input signal, the same delay must be used for each time series. Using different delays would have the same effect as sampling randomly through the time series for each embedded vector in a univariate embedding.

In order to exclude spatially correlated points that are also correlated in time, we examined the drive signal for the length of an average oscillation. We found this to be 30 time steps. We then employed a Theiler (see [13]) window of length 30 time steps in the search for nearest-neighbor points in the continuity test.

The data were normalized and de-meant before embedding. In order to facilitate range searches in the 16-dimensional space, we adapted a kd-tree range search algorithm (see e.g. [2, 3]) to the multivariate embedding. This allowed us to perform each continuity test in about 3 minutes for 100 representative points on the attractor using a 265 MHz G3 processor with 384 MB of memory.

With the 10 data sets for the circuit at damage level 0, we performed 20 source-target combinations of the continuity test in order to determine ϵ_0 and distribution of ϵ^* values between attractors reconstructed from an undamaged circuit's output. We then performed

the continuity test between an undamaged circuit's output as the source and a damaged circuit's output as the target to obtain the average ϵ^* and distribution of ϵ^* for 10 source-target pairs with each level of damage.

In order to assure that the continuity test is a two-way test, we also performed the test for damaged-to-undamaged source-target pairs. These results were similar to those obtained for the undamaged-to-damaged and appear in figure 10.

VI. RESULTS

We summarize the results in figures 7-9. Figure 7 shows the average ϵ^* is larger when there is 2% damage than the average ϵ^* for no damage. We see that the lower end of the 95% confidence interval for the 2% damage scenario is close to the average ϵ^* for the undamaged scenario. The 4% damage shows a similar increase in ϵ^* over that of the undamaged scenario. The 95% confidence intervals for the damage scenarios are also larger than that from the undamaged scenario; in the damage of 6% or more, the interval has more than tripled in size. The evidence of damage is even more striking when we consider the probability distributions seen in figures 8 and 9. The distributions are drawn from 10 runs of the continuity test, each run using a different combination of output from the undamaged circuit and the damaged circuit. We see spreading of the entire probability distribution of ϵ^* as well as clear movement in the location of the peak of the distribution with increasing damage.

We note that although there is a clear difference between the probability distribution for the undamaged and the 2% damage, there is not a large difference between those of the 2% and 4% damage and between 6% and 8% damage. Although the resistors only have a (1%) accuracy in their labelling, the discrepancy does not completely explain this phenomenon. More investigation into both the closeness of the distributions of the 2% and 4% and of the 6% and 8% as well as the apparent jump between 4% and 6% damage is warranted.

It is noteworthy that there is a clearly identifiable change in the nature of the function Ψ even with only 2% change in the circuit. Damage in the 10% scenario was overwhelmingly obvious - we saw a large change in the peak and spread of the probability distribution of the ϵ^* values for the undamaged-to-undamaged Ψ vs. the undamaged-to-damaged Ψ . This points to a loss of continuity/differentiability. We also saw incremental (although not linear) change in the nature of Ψ corresponding to incremental changes in the circuit. Note that

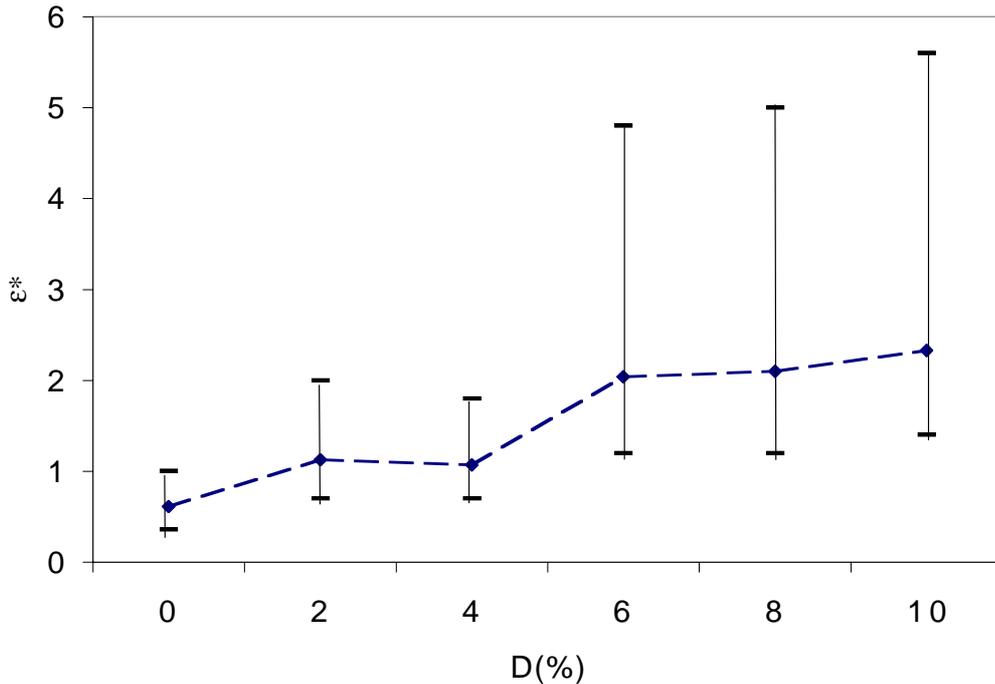


FIG. 7: Changes in average ϵ^* with damage. Maps from undamaged to damaged reconstructions.

by the 95% confidence intervals on figure 7 that the continuity test did *not* indicate *any* damage existed in any of the undamaged circuits. This was reflected in tight distributions with small variance in ϵ^* along with small ϵ^* . Variances for the ϵ^* in the continuity test from undamaged to damaged are listed in table 11.

VII. CONCLUSION

This method of damage detection was extremely sensitive to damage, while at the same time giving consistent results when no damage was present. We saw an incremental change in our statistic with changes in damage, indicating that the method may possibly be used for prognostics as well as diagnostics.

The damage in this structure was confined to one location. Thus, no localization study was possible. A study is currently in progress in which we analyze data from circuits that have the same level of damage in the coupling between different oscillators.

One advantage to this method over current vibrational methods such as described in [6] is that no damage- or structure-specific model is necessary to use this test. The data provide

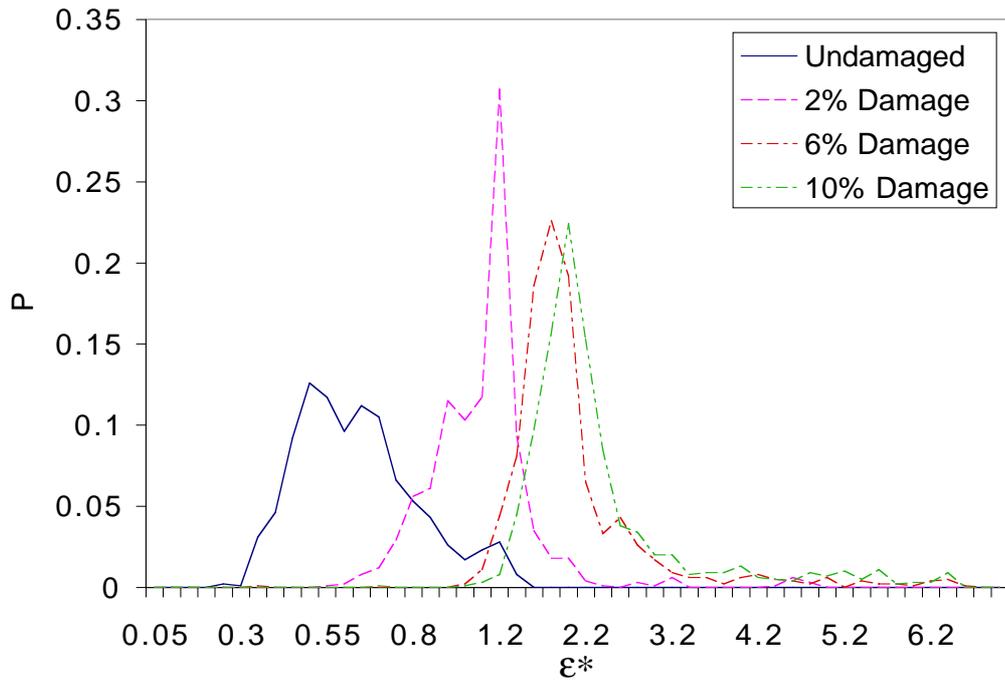


FIG. 8: Distributions of ϵ^* in the Continuity Test. Maps from undamaged to damaged reconstructions.

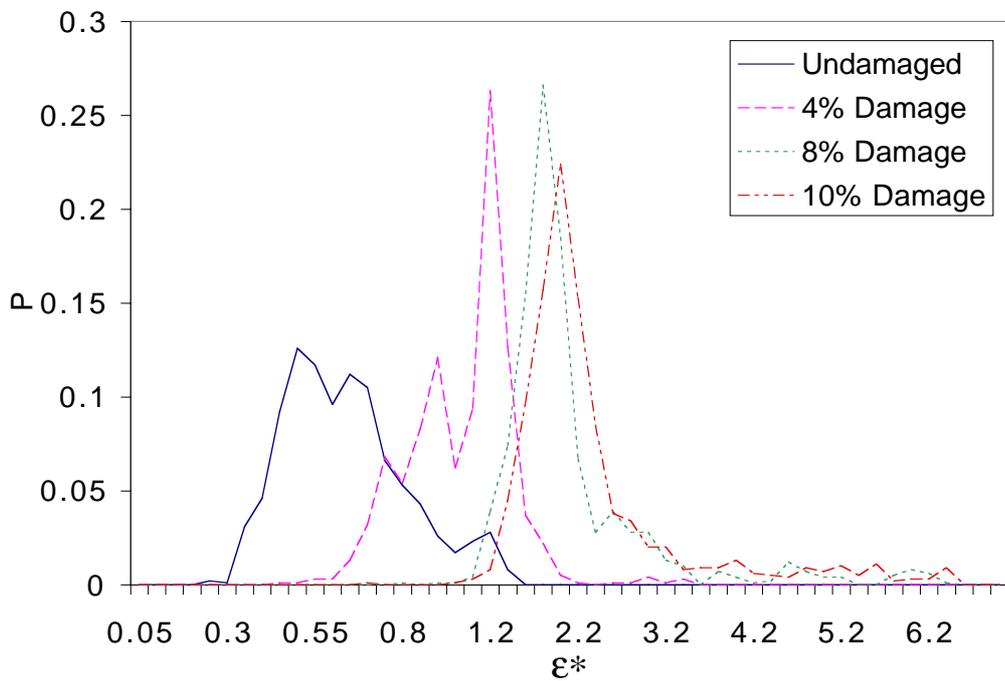


FIG. 9: Distributions of ϵ^* in the Continuity Test. Maps from undamaged to damaged reconstructions.

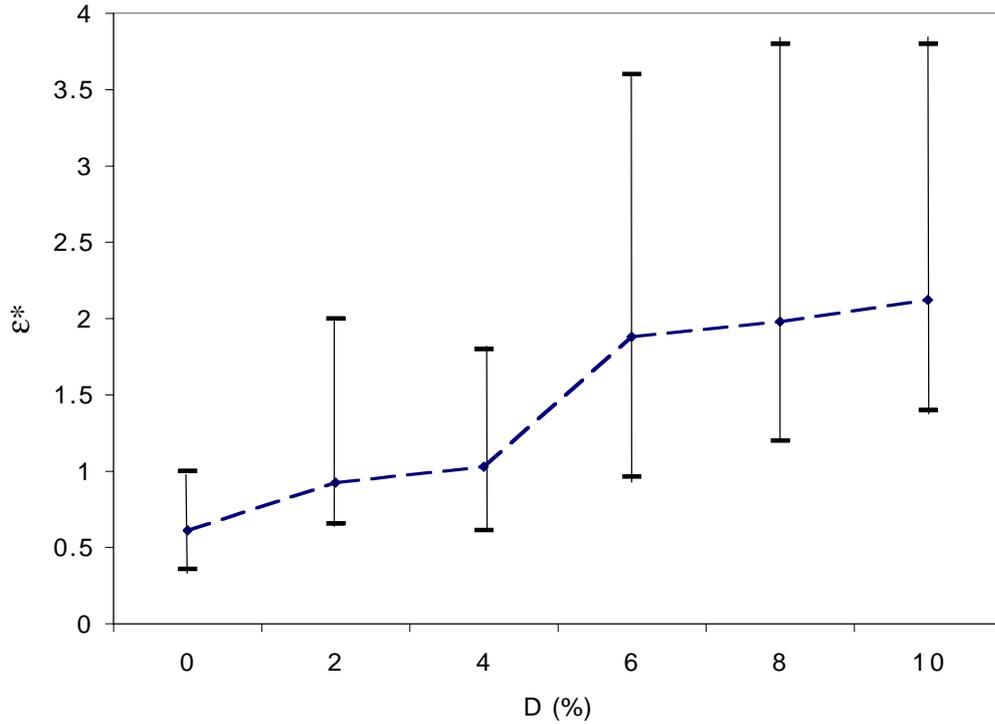


FIG. 10: Changes in average ϵ^* in the Continuity Test for Maps from damaged to undamaged reconstructions

Damage Level	Variance
0	< .05
2%	$\approx .23$
4%	$\approx .12$
6%	$\approx .70$
8%	$\approx .80$
10%	$\approx .90$

FIG. 11: Variance in ϵ^* for various damage levels. Mapping between undamaged and damaged circuit at indicated level.

the model. To arrive at the baseline ϵ^* , one interrogates a pristine structure. Changes in the distribution of ϵ^* indicate change in the stiffness of the structure. The particular damage mechanism does not need to be known in order to detect damage. Using data from multiple sensors, it is possible to localize damage as in [15]; it also may be possible to characterize

various kinds of damage by looking at the maps between attractors reconstructed from individual sensors' data.

- [1] R. Badii, G. Broggi, B. Derighetti, M. Ravani, C. Ciliberto, A. Politi, and M.A. Rubio. Dimension increase in filtered chaotic signals. *Physical Review Letters*, 60(11):979–981, 1988.
- [2] JL Bentley. Multidimensional binary search trees used for associate searching. *Communications of the ACM*, 18(9), 1975.
- [3] JL Bentley and JH Friedman. Data structures for range searching. *Computing Surveys*, 11(4), 1979.
- [4] S. Boccaletti, D.L. Valladares, L.M. Pecora, H. P. Geffert, and T. L. Carroll. Reconstructing embedding spaces of coupled dynamical systems from multivariate data. *Physical Review E*, 65, 2002.
- [5] M.E. Davies and K.M. Campbell. Linear recursive filters and nonlinear dynamics. *Nonlinearity*, 9:487–499, 1996.
- [6] J.E. Doherty. Nondestructive evaluation. In A.S. Kobayashi, editor, *Handbook on Experimental Mechanics*, chapter 12. Society for Experimental Mechanics, Englewood Cliffs, New Jersey, 1987.
- [7] J.M.Nichols, M.D.Todd, M.Seaver, and L.N.Virgin. Use of chaotic excitation and attractor property analysis in structural health monitoring. *Phys Rev E*, 67(016209), 2003.
- [8] M. Kennel, R. Brown, and H. Abarbanel. Determining embedding dimension for phase-space reconstruction using a geometrical construction. In Ott, Sauer, and Yorke, editors, *Coping with Chaos*. Wiley, 1994.
- [9] L.M. Pecora and T.L. Carroll. Discontinuous and nondifferentiable functions and dimension increase by filtering chaotic data. *Chaos*, 6:432–439, 1996.
- [10] L.M. Pecora, T.L. Carroll, and J.F. Heagy. Statistics for continuity and differentiability: An application to attractor reconstruction from time series. *Fields Institute Communications*, 11, 1997.
- [11] Tim Sauer, James A. Yorke, and Martin Casdagli. Embedology. *Journal of Statistical Physics*, 65(3/4):579–616, 1991.
- [12] J. Stark, D.S. Broomhead, M.E. Davies, and J. Huke. Takens embedding theorems for forced

- and stochastic systems. In *Nonlinear analysis, theory, methods and applications*, volume 30, pages 5303–5314. Proc. 2nd World Congress of Nonlinear Analysts, Elsevier, 1997.
- [13] J. Theiler. Spurious dimension from correlation algorithms applied to limited time-series data. *Physical Review A*, 34(12):2427–2431, 1986.
- [14] M.D. Todd, J.M. Nichols, L.M. Pecora, and L.N. Virgin. Vibration-based damage assessment utilizing state-space geometry changes: Local attractor variance ratio. *Smart Materials and Structures*, 10:1000–1008, 2001.
- [15] M.D. Todd, J. Wait, J.M. Nichols, and S.T. Trickey. Joint damage assessment using output-only chaotic attractor property analysis. In *Proceedings of IMAC XXI: A Conference on Structural Dynamics*. Society for Experimental Mechanics, 2003.