Nonlinear-time-series analysis of chaotic laser dynamics

Henry D. I. Abarbanel*
Department of Physics and Marine Physical Laboratory, Scripps Institution of Oceanography, University of California, San Diego, Mail Code 0402, La Jolla, California 92093-0402

Z. Gills
School of Physics, Georgia Institute of Technology, Atlanta, Georgia 30332-0430

C. Liu†
Department of Physics and Institute for Nonlinear Science, University of California, San Diego, La Jolla, California 92093-0402

R. Roy
School of Physics, Georgia Institute of Technology, Atlanta, Georgia 30332-0430
(Received 27 January 1995; revised manuscript received 14 September 1995)

We report on the analysis of experiments on a neodymium-doped yttrium aluminum garnet laser with an intracavity frequency-doubling crystal. Three modes of the laser were excited in differing polarization configurations. The total intensity of infrared light was observed and then analyzed using methods of nonlinear-time-series analysis. We present clear evidence using global false nearest neighbors that when all polarizations are parallel, the intensity is chaotic with two positive Lyapunov exponents and the system can be embedded in dimension 7. The noise level in this operating condition, which we call type I chaos, is small. When one of the polarizations is perpendicular to the others, the intensity is again chaotic with positive Lyapunov exponents, but there is substantial noise in the signal of high dimensional origin, and no finite embedding dimension appears possible. We call this type II chaos. We suggest that the origin of this phenomenon is the intrinsic quantum noise associated with the generation of green light, which is 25 times more intense in the type II operating configuration than in the first. In past experiments with this system we have found that the type I chaos can be controlled to unstable periodic orbits while type II cannot. In each type of chaotic laser operation we use local false nearest neighbors to demonstrate that the local dimension of the dynamics is 7. This means seven differential equations can capture the full dynamics of these regimes of the laser. We evaluate local and global false nearest neighbors to support our conclusions and determine the Lyapunov spectrum of each type of chaotic behavior. The predictability of type II chaos is shown to be much less than that of type I, and we make local polynomial models in reconstructed-state space to demonstrate that we can predict rather well for type I chaos. Finally we suggest a fairly standard model for the interaction of the infrared light with the nonlinear frequency doubling medium and with a two-level of the active medium.

PACS number(s): 42.60.V, 422.50.Ne, 42.50.Lc, 05.45.+b

I. INTRODUCTION

In the operation of a neodymium-doped yttrium aluminum garnet (Nd:YAG) laser with an intracavity potassium titanyl phosphate (KTP) crystal, irregular fluctuations of the total output intensity in infrared light ($\lambda \approx 1.064\ \mu$m) are commonly observed. These fluctuations are chaotic as they have at least one positive Lyapunov exponent associated with their evolution [1,2], and they have a broad, continuous Fourier power spectrum with a peak near 60 kHz. The observation and characterization of this chaotic signal are now somewhat familiar, although we bring tools to the analysis of the chaos that have not previously been applied to this physical problem. The focus in this paper is on aspects of this chaotic laser that are tied to the quantum mechanical generation of green light ($\lambda \approx 0.532\ \mu$m) via the KTP crystal. The green light leaves the cavity as one of the mirrors is transparent at its wavelength.

Operating with three active cavity modes for the infrared light, we observed two distinctly different irregular time series depending on the polarizations of the light. When the modes were all polarized parallel to each other, the chaotic oscillations of the total infrared intensity showed clear low-dimensional behavior. This class of oscillations, which we term type I chaos, was also accompanied by a very low level of green light. In this setup we anticipate that the dynamical equations governing the intensity of the infrared light and the gain in the active medium would be quite accurate in their semiclassical form.

We also observed a second kind of chaotic motion, type II chaos, which occurs when one of the modes of infrared light is polarized perpendicular to the other two. In this case the production of green light is very strong, and the false nearest-neighbor statistic that determines the integer dimension in which the dynamics can be captured shows that this dimension is not small. The ''noise'' seen by this statistic we

*Also at the Institute for Nonlinear Science, University of California, San Diego, La Jolla, CA 92093-0402. Electronic address: hdia@hamilton.ucsd.edu

†Electronic address: clif@rutherford.ucsd.edu

associated with the intrinsic fluctuations accompanying the generation of the green light. For type II chaos because the apparent noise levels are higher we anticipate that the dynamical equations may be semiclassical but significantly influenced by fluctuation terms.

In this paper we review these observations beginning in the next section with a discussion of the experiment and the data collection. In Sec. III we discuss the average mutual information and false nearest-neighbor methods—local and global—we use to analyze the data, and draw some of the implications just suggested. In particular we use global false nearest neighbors to demonstrate that the global dimension d_E in which the attractor of type I chaos is unfolded is $d_E=7$. Type II chaos has enough noise that no global embedding dimension can be extracted. For each type of chaos we are able to deduce the local dynamical dimension of $d_L=7$ using local false nearest neighbors. In this section we also discuss the evaluation of Lyapunov exponents for these chaotic data. Interestingly we find two positive Lyapunov exponents for type I chaotic behavior, and three positive exponents for type II. Using various techniques for controlling chaos, we have been able to control type I chaos to periodic behavior, but have been unsuccessful with type II chaos [3,4]. We attribute this to the additional positive Lyapunov exponents as well as the higher intrinsic noise level associated with more green light production.

In Sec. IV we discuss making models of the evolution of the laser system in a reconstructed phase space made out of the observed total infrared intensity and its time delays. These models are ''black box'' models that do not rely on knowing any of the physics of the laser system. They are interpolation rules in the reconstructed phase space that allow us to predict from a new data point, here a given total infrared intensity, what the future evolution of that intensity will be. The method, described in some detail below, uses knowledge of how one evolves in time from one neighborhood of a point on the orbit to the next neighborhood in phase space.

In Sec. V we use a familiar model of the interaction of the infrared light with the nonlinear susceptibility of the KTP crystal to produce green light. This model is expanded to describe our multimode finite-length laser. The numerical work to verify this model is still in progress.

Our final section is a summary of all this and includes some suggestions for further investigation of these phenomena.

II. EXPERIMENTAL SETUP AND DATA PREPARATION

A. Experimental setup

The basic elements of the laser system are a Nd:YAG crystal pumped by a diode laser that is in the same cavity as a KTP crystal. The nonlinear KTP crystal serves as the frequency-doubling element. One end of the laser cavity is formed by the high reflection coated flat facet of the Nd:YAG crystal. This facet is highly reflecting at both the infrared fundamental (1064 nm) and at the doubled green (532 nm) wavelength and is highly transmissive at the pump wavelength (810 nm). The back end of the laser cavity is formed by a curved output coupler that is highly transmissive for the green light and highly reflecting for infrared light [1]. The KTP crystal is antireflection coated for both the fundamental and doubled wavelengths. The green light is thus not resonant in the cavity and acts as a nonlinear loss mechanism.

The laser system can display steady-state, periodic, quasiperiodic, and chaotic intensity fluctuations when operated with three or more longitudinal modes. The characteristic time scale on which these fluctuations occur is that of the relaxation oscillations, which are inherent to the laser system. They are the result of energy exchange between atoms in the lasing medium and light in the laser cavity. Without the KTP crystal in the cavity these oscillations are normally heavily damped and stabilized by this damping. The doubling process provides the nonlinear loss mechanism that destabilizes these oscillations. For the data presented in this paper, the relaxation oscillation frequency is approximately 60 kHz.

The mode structure of the YAG laser output was monitored using a confocal Fabry-Pérot interferometer with a free spectral range of 8 GHz. Each of the observed longitudinal modes was polarized along one of two orthogonal directions. The relative polarization direction of each mode was determined by separating the laser output using a polarizing beam-splitter cube. The polarized output was monitored along one direction and then the cube was rotated to monitor the output along the orthogonal polarization direction. The total intensity of infrared light was monitored with a photodiode that converted the laser intensity to voltages that were viewed using a digital oscilloscope.

In the present experiments the system parameters were adjusted to obtain chaotic behavior in laser operation with three infrared cavity modes. Two distinct polarization configurations were selected. An appropriate orientation of the crystal axes allowed us to select these configurations. The linear cavity loss and the pump level, set to about twice the threshold pump power, were similar for the two configurations. The total intensity in the infrared $I(t)$ was observed with a photodiode having a rise time of less than 1 ns and was sampled using a 100-MHz eight-bit digital oscilloscope capable of storing 10^6 samples. In Fig. 1 we show $I(t)$ when all three modes were polarized parallel to each other; we call this type I chaos. In Fig. 4 we show $I(t)$ with one mode polarized perpendicular to the other two: we call this type II chaos.

Even in the time traces we can see the distinction between these two operating regimes. Type I consists of long ''bursts'' of relaxation oscillations, while type II appears far more irregular. During type I operation very little green light, less than 1 μW, was observed, while more than 25 μW of power in green light accompanied type II activity. This is consistent with the linear stability properties of the macroscope equations we will present later. If all three modes are parallel polarized as in type I behavior, the laser can become unstable with a very small coupling in the KTP crystal, but very little green light is produced. If one of the modes is polarized perpendicular to the other two, a very small value of the coupling results in appreciable sum frequency generation, instability to chaotic operation, and roughly two orders of magnitude larger intensity of green light than in type I operation.

B. Data preparation

The resolution of the digital sampling oscilloscope (eight bits or 0.4%) is too low for some of the data analysis, notably the Lyapunov exponent calculation. In order to improve the resolution of the data sets and still have the large number of points needed to analyze data in high dimensions, the oscilloscope is set to sample the data at $f_s = 10$ MHz and the entire 10^6 sample memory is used. This rate is 200 times higher than the 60-kHz relaxation oscillations of the dynamics ensuring that the detailed evolution of the signal is captured and aliasing does not occur. Recall that the sampling rate f_s must be at least twice the highest frequency found in the signal to prevent aliasing.

For the calculations where resolution was not critical, the data sets are downsampled by a factor of eight resulting in an effective sampling rate of 1.25 MHz or a sampling period of 800 ns. This rate is still 20 times higher than the relaxation oscillation rate, and so the first minimum of the mutual information is found at four or five samples or one-quarter of the relaxation oscillation period. Note that in downsampling, the broadband noise level (due to high-dimensional dynamics) is neither increased or decreased. This is important when one is trying to determine the noise level using the global false neighbors algorithm.

When the eight-bit resolution is insufficient, the raw data sets were interpolated using a digital linear filter. This filter is designed to remove frequencies from 500 kHz to the Nyquist frequency $f_s/2 = 5$ MHz and pass all frequencies below 500 kHz. Since the signal due to the dynamics alone has no frequency as high as 500 kHz, no dynamical information is lost.

On the other hand, the quantization noise, which is assumed to be white (equally distributed over all frequencies) up to the Nyquist frequency, is cut down by 90%. Thus, by reducing the quantization noise or error by a factor of 10, we have gained slightly more than three bits of resolution.

The digital filter is implemented by performing a discrete convolution of the data with the impulse response, which is the inverse Fourier transform of the filter frequency response. A Hamming window is applied to the impulse response to reduce frequency anomalies at the cutoff frequency (Gibb's phenomenon) [5]. This impulse response is chosen to be symmetric in time so the filter exhibits linear phase; i.e., it will cause no relative phase shifts at different frequencies.

After performing the digital convolution, the data are downsampled from the original 10 MHz to 1.25 MHz. Since the filter has already removed frequencies above 500 kHz, no aliasing occurs and the data have 11 bits of resolution.

III. NONLINEAR ANALYSIS OF THE DATA

In the analysis of the chaotic laser data we employ several ideas that are covered in detail in Ref. [2], but that we will discuss here for completeness. The primary goal of our analysis is to establish a d_E-dimensional space of data vectors whose components are the measured total intensity of observed infrared light $I(t)$ and its time delays. d_E is an integer. If the infrared light is sampled every τ_s in time, then we wish to form the d_E-dimensional data vectors

$$\mathbf{y}(n) = [I(n), I(n+T), I(n+2T), \ldots, I(n+(d_E-1)T)], \tag{1}$$

where $I(n) = I(t_0 + n\tau_s)$ and T is an integer multiple of the sampling time τ_s. We need a method for choosing the time delay T and the required integer embedding dimension d_E. These vectors $\mathbf{y}(n)$ then constitute our data and provide us with the reconstructed phase space of the dynamical system. It is in this d_E-dimensional space that we evaluate all the relevant dynamical quantities of the laser.

A. Mutual information

To choose the time delay T we utilize the nonlinear ''correlation'' function of average mutual information. This answers the question: how much, in bits, does one learn on the average about a set of measurements $A = \{a_j\}$ from a set of measurements $B = \{b_n\}$. Here the two sets of measurements are the set of intensities $I(n)$ at times $n = 1,2,\ldots,N$ and the set of intensities $I(n+T)$ at times $n = 1,2,\ldots,N$. The mutual information between these two sets of measurements is

$$\log_2\left[\frac{P(a_j, b_n)}{P(a_j)P(b_n)}\right], \tag{2}$$

where $P(\cdot,\cdot)$ is the joint probability distribution for the two measurements $I(n)$ and $I(n+T)$. The quantity $P(\cdot)$ is the individual probability distribution for either of the measurements. The correlation function, which is the *average mutual information* over all measurements, is

$$\mathscr{I}(T) = \sum_{a_j, b_n} P_{AB}(a_j, b_n) \log_2\left[\frac{P_{AB}(a_j, b_n)}{P_A(a_j)P_B(b_n)}\right], \tag{3}$$

where the a_j are the $I(n)$ and the b_n are the $I(n+T)$. The quantity $\mathscr{I}(T)$ gives us a quantitative, nonlinear measure of the independence of the measurements $I(n)$ and $I(n+T)$. If these are completely independent in a nonlinear fashion, then $P_{AB}(a_j, b_n) = P_A(a_j)P_B(b_n)$, and the mutual information is zero. If $T = 0$, then the average mutual information is just the entropy of the set of measurements. We seek a value of T at which the measurements are somewhat independent, but not totally independent. In the latter case they would not serve well as coordinates in the d_E-dimensional state space of the vectors $\mathbf{y}(n)$. If T were quite large the measurements would be independent, since the source of the signal is chaotic and thus unstable everywhere in state space. This means that for large T the measurements $I(n)$ and $I(n+T)$ are essentially random with respect to each other. For T too small, not enough time has passed for the measurements to be somewhat independent of each other. We utilize the *prescription* of choosing that time lag T, which is the first minimum of the average mutual information [6,2] with which to construct our vectors $\mathbf{y}(n)$. This prescription gives us an intermediate value of T to use in forming the vectors $\mathbf{y}(n)$.

For data of type I chaos, we display in Fig. 1 a sample time series from our observations. The original sampling time for these data was 100 ns. We then downsampled this data by a factor of eight because the original data were oversampled; this was done after the preparation of the data as described in Sec. II B above. The effective sampling time for all the data we consider in this paper is then $\tau_s = 800$ ns.

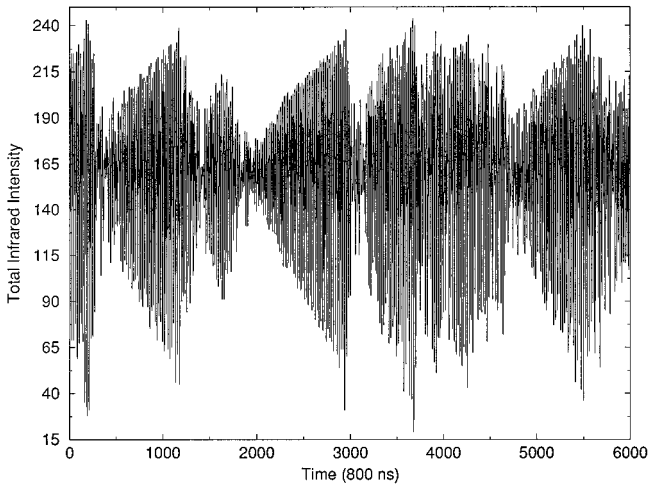

FIG. 1. A typical time trace for type I chaos. The sampling time is $\tau_s = 800$ ns. 6000 points out of the 125 000 observed are displayed. The laser was operating with three infrared cavity modes all polarized parallel to each other.

Figure 1 shows 6000 data points out of the 125 000 collected. The Fourier power spectrum of these data is in Fig. 2, and we can see broad spectral features near 60 kHz and apparent harmonics of that frequency. The broadband nature of the spectrum is characteristic of chaotic motion. Figure 3 shows the average mutual information evaluated from these data. We see that the average mutual information has its first minimum near $T\tau_s = 5\tau_s = 4$ μs. This was the location of the first minimum in $\mathscr{I}(T)$ for each of the type I data sets we examined.

For data from a type II chaos trace, the time series is shown in Fig. 4 and the Fourier spectrum in Fig. 5. $\tau_s = 800$ ns again. The spectral features seen in type I chaos are washed out with other spectral ''peaks'' visible in these data. Once again we evaluated the average mutual information and see a typical result in Fig. 6 from one of our data sets we examined. The first minimum for $\mathscr{I}(T)$ is again near $T = 5$ for each of the type II data sets. Here it was at $T = 4$.

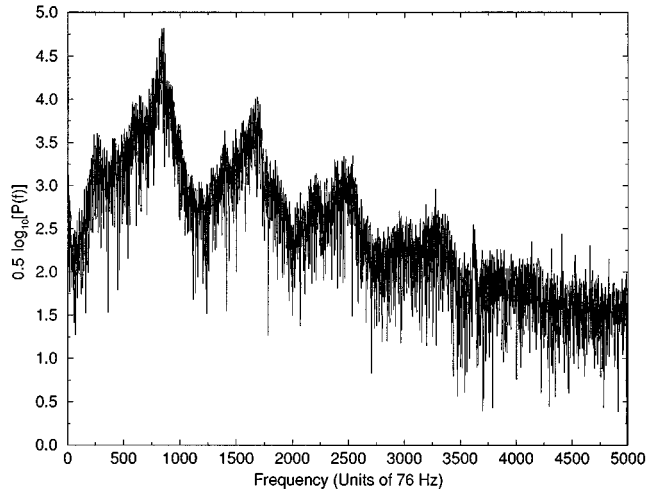

FIG. 2. The Fourier power spectrum of the type I chaos seen in Fig. 1. There is a broad peak in the spectrum near 60 kHz and along with other broad spectral features.

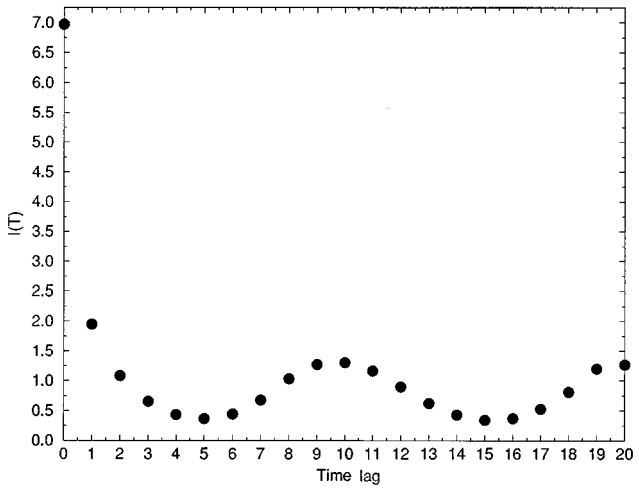

FIG. 3. The average mutual information $\mathscr{I}(T)$ for type I chaos. 125 000 samples were used. A clear first minimum at $T = 5$ corresponding to 4 μs is seen.

The distinction between the type I and the type II data comes when we examine another characteristic of the signals: the false nearest neighbors.

It is important to restate that this choice, or any choice, of T is a prescription. Mutual information should give a value for T that is indicative of the relevant times for nonlinear processes to manifest themselves, but a cautious consumer of the methods we are using will investigate all results for a selection of T values around the first minimum of the average mutual information. We have done so in each of the computations we report as we go along and found the precise value of T not to be important. $T \pm 1$, say, would do as well.

B. False nearest neighbors

In order to determine the required embedding dimension d_E to use in the data vectors $\mathbf{y}(n)$ we employ a statistic about

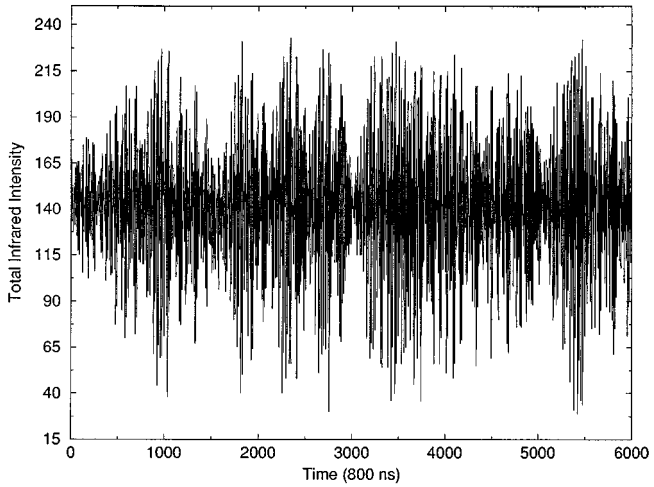

FIG. 4. A typical time trace for type II chaos. The sampling time is $\tau_s = 800$ ns. 6000 points out of the 125 000 observed are displayed. The laser was operating with three infrared cavity modes where two are polarized parallel to each other and the other is polarized perpendicular to the first two.

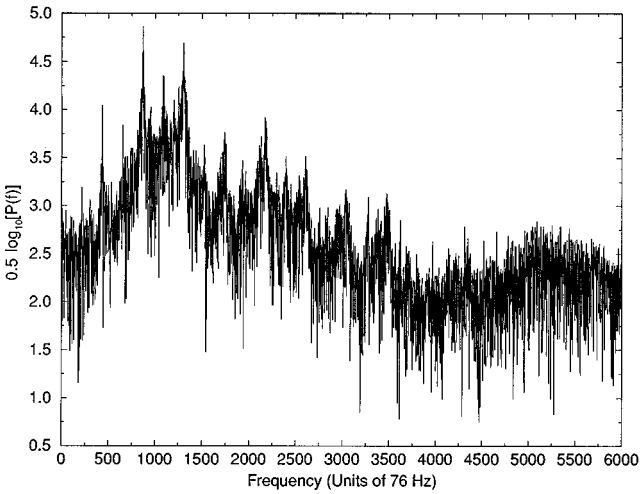

FIG. 5. The Fourier power spectrum of the type II chaos seen in Fig. 4. The broad peak near 60 kHz seen in type I chaos has been further broadened.

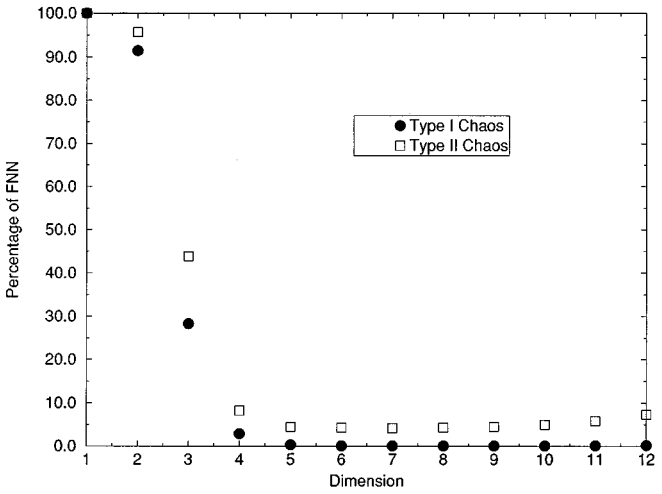

FIG. 7. The percentage of global false nearest neighbors for both type I (solid symbols) and type II chaos (open symbols). For type I chaos we used $T=5$ for the state space reconstruction while for type II, $T=4$. Both operating conditions show some noise, but the levels of residual false neighbors is much higher in type II chaos where the green production is substantially larger.

these data known as false nearest neighbors [7]. The construction of the d_E-dimensional data vectors is purely geometric once a time delay T has been chosen. One is seeking a multivariate space in which the orbits $\mathbf{y}(n)$ do not overlap because of projection from a higher dimension. The main idea is that if the data come from a system governed by autonomous differential equations or discrete time maps with time-independent parameters, then orbits cannot cross each other. If they appear to cross, it is because we are viewing them in a dimension that is too small or that is smaller than the number of active degrees of freedom in the source of our observations.

To establish whether the space has become big enough we look at each point $\mathbf{y}(n)$ and its nearest neighbor $\mathbf{y}^{\mathrm{NN}}(n)$ in dimension d. The time index k associated with $\mathbf{y}^{\mathrm{NN}}(n)=\mathbf{y}(k)$ need bear little resemblance to the time index n of the data point we are examining. As data fold back on themselves in dissipative systems, the neighborhood of $\mathbf{y}(n)$ may be populated by points such as $\mathbf{y}^{\mathrm{NN}}(n)$ of quite different time index. Now we ask what happens to the distance between $\mathbf{y}(n)$ as seen in dimension $d+1$ where it is the vector

$$\mathbf{y}(n)=(I(n),I(n+T),\ldots,I(n+Td)), \tag{4}$$

and the vector $\mathbf{y}^{\mathrm{NN}}(n)$ in dimension $d+1$, where it becomes

$$\mathbf{y}^{\mathrm{NN}}(n)=\mathbf{y}(k)=(I(k),I(k+T),\ldots,I(k+Td)). \tag{5}$$

If the distance in dimension $d+1$ is large, we label $\mathbf{y}^{\mathrm{NN}}(n)$ a *false nearest neighbor* of $\mathbf{y}(n)$. When d becomes large enough, the number of false neighbors will go to zero. At that dimension the attractor is unfolded in the coordinate system defined by our choice of time delay T. Now the idea of ''large'' distances requires some threshold for deciding that a nearest neighbor $\mathbf{y}^{\mathrm{NN}}$ is false. The choice of false neighbor turns out to be independent of this threshold over a wide range of choices for its value [7].

For our type I data the percentage of false nearest neighbors is shown in Fig. 7 for one of our data sets. The other data sets produce quite similar results. The interpretation of this is quite clear: the attractor associated with type I chaos is captured in a low-dimensional space with $d_E\approx 5$. Our analysis of type II chaos yields quite a different result. We show in Fig. 8 the false nearest neighbors for $4\leqslant d_E\leqslant 8$ for data sets of type I and type II. We see that the percentage of global false nearest neighbors for type I chaos falls to zero at $d_E=7$. The percentage of false nearest neighbors does not fall to zero for any dimension for the type II data, and this indicates the presence of a high-dimensional noise in these data. We note that when the chaos is type II, the infrared light is accompanied by a large amount of green light associated with the conversion of infrared photons in the KTP crystal.

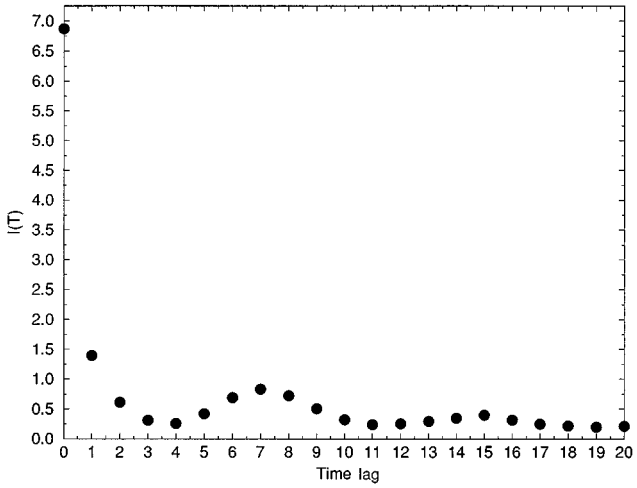

FIG. 6. The average mutual information $\mathscr{I}(T)$ for type II chaos. 125 000 samples were used. A clear first minimum at $T=4$ corresponding to 3.2 μs is seen.

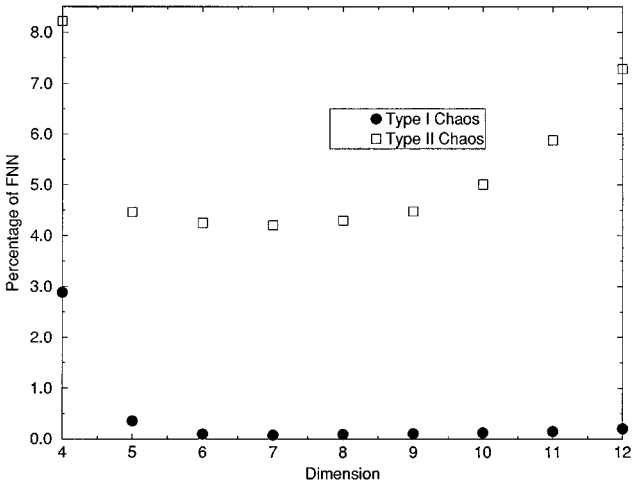

FIG. 8. An enlargement of the global false neighbors statistic to emphasize the residual level of false neighbors for type II chaos.

C. Local false nearest neighbors

The integer dimension d_E evaluated above tells us the global dimension required to unfold the attractor from its projection on the observation axis $I(n)$. The local dimension d_L of the dynamics may be less than or equal to d_E. If $d_E > d_L$, it means that the particular coordinate system, namely, that of the time delay vectors $\mathbf{y}(n)$, twists it about so that unfolding the attractor from the observations requires additional dimensions to "untwist it." The local or dynamical dimension is the same in any coordinate system and represents the integer number of differential equations required to describe the local evolution of the system.

We use the method of *local* false nearest neighbors [8] that starts with a working dimension $d_W \geq d_E$ with d_E as determined by the global false-neighbors algorithm described above. All distances between points $\mathbf{y}(n)$ are computed in d_W. Now we choose the local dimension d starting with $d=1$ and proceeding to $d=d_W$ asking at each dimension how well we can predict the evolution of a cluster of neighbors about $\mathbf{y}(n)$ for each $n=1,2,\ldots,N$ on the attractor. In a dimension d, which is lower than the true local dimension d_L, there will be points in a neighborhood which got there by projection and not by a dynamical rule. The ability to predict where these points go will be very bad. As we increase d up to d_L our ability to predict will improve until it levels off and becomes independent of the estimated local dimension d as well as the number of neighbors we use to define a neighborhood whose evolution we predict. In dimensions $d_L \leq d \leq d_W$ predictability should become independent of d as we have enough dynamical degrees of freedom already to capture the evolution of the data. The quality of prediction is determined by how far ahead we wish to predict, and by how large an error we tolerate until we say the prediction has failed.

To be more precise about the details of the method of local false nearest neighbors we note that what is involved in the prediction is a local map from one neighborhood of the attractor to another neighborhood of the attractor. The members of a neighborhood are determined by distances evaluated in a dimension $d_W \geq d_E$ with d_E the global embedding dimension determined by the false-nearest-neighbors method. This assures that all neighbors are true neighbors. Then in a neighborhood, we ask if a local model in dimension $d \leq d_W$ accurately relating the neighbors of the point $\mathbf{y}(n)$ to the neighbors of $\mathbf{y}(n+1)$ can be made. This local model is a rule that gives $\mathbf{y}(n+1)$ in terms of a polynomial constructed out of the vectors $\mathbf{y}(n)$. The coefficients in this polynomial rule are determined by a least-squares fit minimizing the residual errors in the rule that takes $\mathbf{y}(n) \rightarrow \mathbf{y}(n+1)$. This rule is used to predict ahead a time T equal to the mutual information time determined before. If, in predicting ahead this amount of time, the error between the prediction and the known observed data point is larger than a certain fraction β of the attractor size, we deem it a bad prediction. The fraction of bad predictions P_K is collected as we move over the whole attractor, and this is displayed against the dimension d of the local model for various choices of the number of neighbors N_B of the points $\mathbf{y}(n)$ on the orbit. We seek a dimension where the model, represented by its fraction of bad predictions, becomes independent of the local dimension and the number of neighbors used to establish the model. The d at which this occurs determines the dimension d_L of a good local model. As one changes the fraction β of the attractor size that defines the allowed error sphere or changes the time that one predicts ahead, the fraction of bad predictions P_K will move up and down in absolute value. The dimension d_L at which P_K becomes independent of d and N_B will remain the same. As it is d_L that one wishes to extract from the local false-nearest-neighbors statistic the absolute size of P_K is of no special significance.

In all evaluations of local false nearest neighbors we report here we chose $d_W=15$ and $\beta=0.37$. The "size" of the attractor was chose as the root-mean-square variation of the intensity values $I(n)$ about their mean over the whole data set. Other natural choices for the extent of the overall attractor in phase space lead to variations in the absolute level of P_K but not to changes in the choice of d_L. In our computations we always chose $N_B=10$, 25, 50, and 75 as illustrative values to assure ourselves that P_K had become independent of N_B as well as of the dimension d. In Fig. 9(a) we show the fraction of bad predictions P_K for type I chaos when we predict forward in time. Clearly P_K becomes independent of the local dimension d and of N_B at $d_L=7$. If we predict backward in time, as shown in Fig. 9(b), we draw the same conclusion. This is consistent with the behavior of type I chaos seen in Fig. 8 where the number of global false nearest neighbors falls to zero at $d_E=7$, then remains there. Local false nearest neighbors are much more sensitive to fine structure on the attractor than global false neighbors. The latter presents a kind of global average over all regions of phase space, so regions that exhibit the highest dimensional structure may occupy only a small percentage of the total phase space. This small percentage of space regains proper importance when local quantities are computed, as we are doing now.

In Figs. 10(a) and 10(b) we show the result of the same calculation for type II chaos. Note that the P_K values are much higher for type II chaos, reflecting the presence of high-dimensional noise in the data. This noise is not of such high amplitude as to ruin completely the possibility of making local predictions, but it certainly enormously erodes the

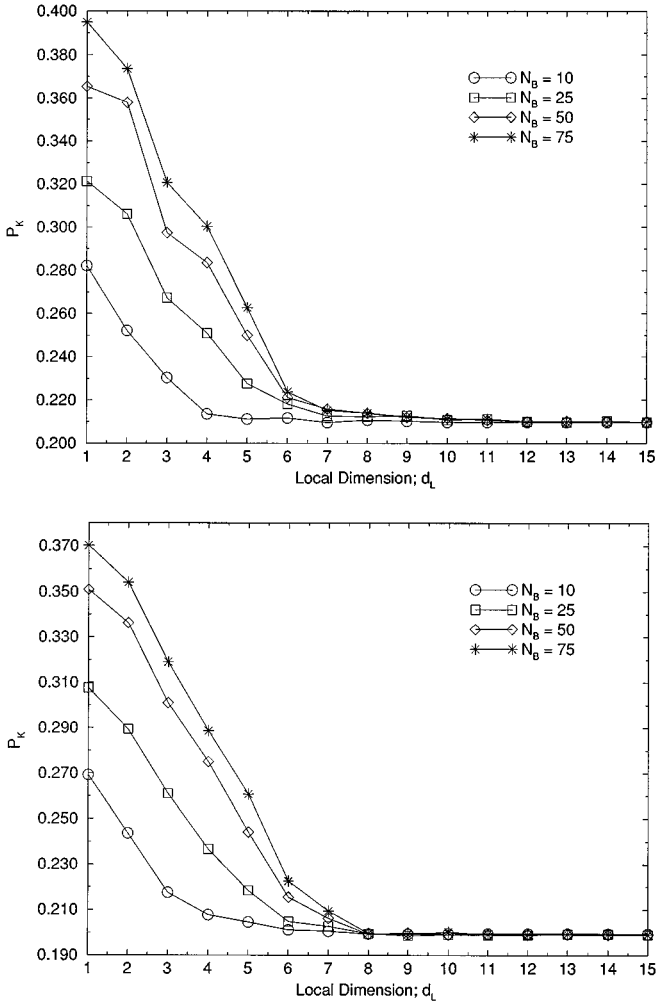

FIG. 9. Local false nearest neighbors for type I chaos using $T=5$ and 120 000 points from the time series. The computation done (a) forward and then (b) backward on the data. There is a clear indication that at $d_E=7$ the predictability of these data has become independent of the number of neighbors and the embedding dimension. β, here set to 0.37, defines the size of the error ball within which a good prediction must fall after T steps forward (or backward) in predicting. The error ball is β times the overall size of the attractor. See [8] for more details on this parameter choice in this algorithm.

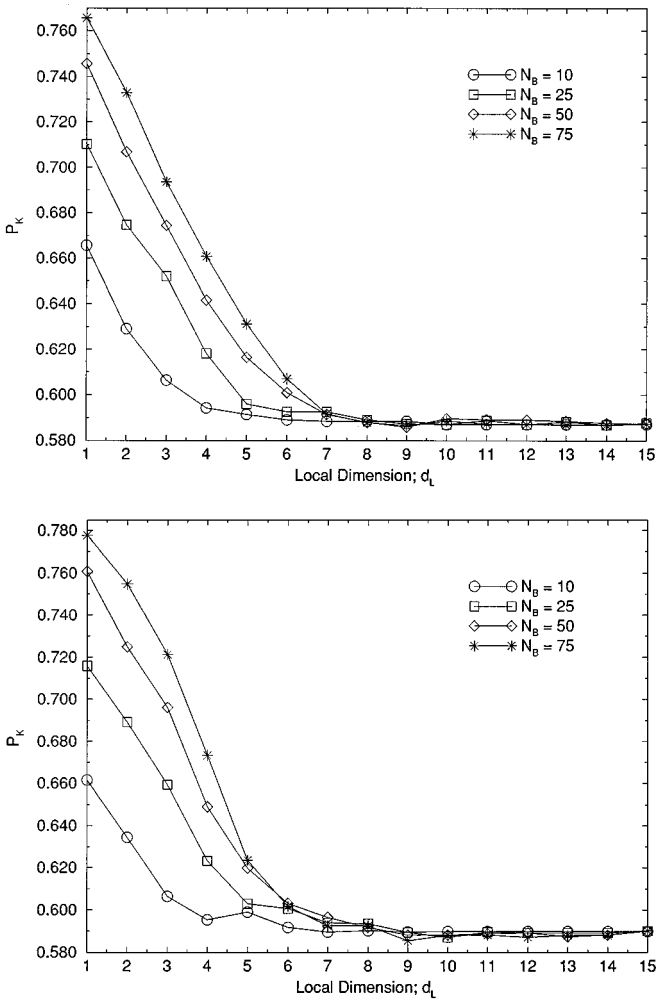

FIG. 10. Local false nearest neighbors for type II chaos using $T=4$ and 120 000 points from the time series. The computation done (a) forward and then (b) backward on the data. There is a clear indication that at $d_E=7$ the predictability of these data has become independent of the number of neighbors and the embedding dimension. Note the much higher percentage of *unpredictable* points on the attractor here compared to type I chaos. The level of unpredictability for $d \geqslant d_L$ is nearly three times that seen for type I chaos. This is a direct result of the higher noise level in type II chaos. Nonetheless the local false-nearest-neighbor statistic is seen to be quite robust against noise. β, here set to 0.37, defines the size of the error ball within which a good prediction must fall after T steps forward (or backward) in predicting. The error ball is β times the overall size of the attractor. See [8] for more details on this parameter choice in this algorithm.

quality of those predictions. The level of bad predictions for the same parameter settings for both types of chaos leads to about 20% bad predictions for type I chaos and nearly 60% bad predictions for type II chaos. At the same time there is a clear indication that the dynamical dimension of the system giving rise to the observations is $d_L=7$ in each case. This is a very nice result in that it shows that low levels of noise do not impede our ability to identify the number of differential equations required to describe the data.

Seven degrees of freedom could well have been anticipated from physical reasoning. We have three modes of the infrared field each of which has a creation and an annihilation operator describing it. So we have six degrees of freedom from the electromagnetic field. The green field is significantly damped by its not being a cavity mode, so associated with that field we should expect some large and negative Lyapunov exponents. Finally we have the atomic degrees of freedom in the active medium. We anticipate a single gain equation associated with the population inversion of the level responsible for the principal transition near 1.064 μm. This will give us the seven degrees of freedom seen in the experimental data.

If we provide more degrees of freedom in the form of differential equations for Heisenberg operators or density matrix elements, we expect to find large, negative Lyapunov exponents associated with the damping of these quantities. The local false nearest neighbors results tells us the number of "active" degrees of freedom out of the many we could

anticipate entering this problem. ''Active degrees of freedom'' is best defined by the example we are discussing now: namely, those dynamical variables that are not substantially removed from the dynamical description of the physical situation by damping or losses. Once again those variables that are suppressed by the losses would show up in exhibiting large, but negative, Lyapunov exponents were we to evaluate them in dimensions larger than that indicated by local false neighbors, here $d_L=7$.

D. Local and global Lyapunov exponents

The evolution of small perturbations to an observed orbit $\mathbf{y}(n)$ is governed by the linearized equations of motion, whatever they may be. As we do not know them from looking at the data alone, we assume there is an underlying nonlinear process $\mathbf{y}(n)\rightarrow\mathbf{y}(n+1)=\mathbf{F}(\mathbf{y}(n))$ that moves the system ahead one sampling time τ_s. A small perturbation $\mathbf{w}(n)$ to the orbit $\mathbf{y}(n)$ satisfies

$$\mathbf{y}(n+1)+\mathbf{w}(n+1)=\mathbf{F}(\mathbf{y}(n)+\mathbf{w}(n)),$$

$$\mathbf{w}(n+1)=\mathbf{DF}(\mathbf{y}(n))\cdot\mathbf{w}(n)+O(w^2), \tag{6}$$

where the Jacobian matrix is

$$\mathbf{DF}(\mathbf{x})_{ab}=\frac{\partial F_a(\mathbf{x})}{\partial x_b}, \tag{7}$$

and $a,b=1,2,\ldots,d_L$.

The eventual growth or shrinkage of the perturbation under this linear evolution rule is determined by the eigenvalues $e^{\lambda_a(\mathbf{x},L)}$ of the Oseledec matrix [2,9]

$$\mathcal{O}(\mathbf{x},L)=[(\mathbf{DF}^L(\mathbf{x}))^T\cdot\mathbf{DF}^L(\mathbf{x})]^{1/2L}. \tag{8}$$

The matrix $\mathbf{DF}^L(\mathbf{x})$ is the composition of L Jacobian matrices along the observed orbit $\mathbf{y}(n)$ starting at location $\mathbf{x}$. As $L\rightarrow\infty$, the $\lambda_a(\mathbf{x},L)\rightarrow\lambda_a$, which are the usual global Lyapunov exponents. The λ_a are independent of $\mathbf{x}$ in the basin of attraction of the attractor. They are invariants of the dynamics and characterize it. They are also independent of the coordinate system in which they are evaluated. To determine the λ_a reliably we need to know the value of d_L as this is the dimension of the dynamics. If we work in a space with $d>d_L$, then $d-d_L$ of the eigenvalues of the Oseledec matrix will be false, and we need a reliable rule to establish which are true and which are not. Similarly if we work in $d<d_L$, we will not have unfolded the local dynamics in such a way that would allow the correct evaluation of the $\mathbf{DF}(\mathbf{x})$, and thus the λ_a would be in error.

Using $d_E=d_L=7$, we have computed the $\lambda_a(\mathbf{x},L)$ for a large number of starting location $\mathbf{x}$ on the attractor, then determined the value of these quantities as a function of the number of steps we look ahead of these starting points. We used 5000 starting points and carried the calculation out 2048 steps ahead of each of these locations. This allows us to define an average local Lyapunov exponent.

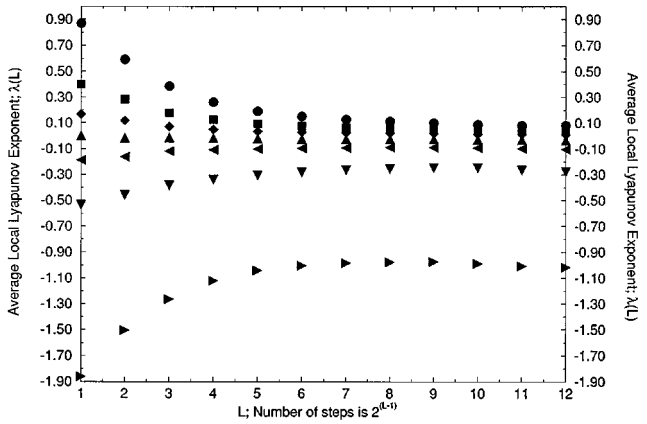

FIG. 11. The average local Lyapunov exponents for our sample of type I chaos. A global and a local dimension of $d_E=d_L=7$ has been used, and 120 000 points from the time series were utilized. There are two positive Lyapunov exponents and one zero exponent indicating that differential equations describe the underlying dynamics. The Lyapunov dimension D_L for this is about 4.95 ± 0.1 telling us that the last large, negative exponent is not very important dynamically. This is consistent with the fall of the global false nearest neighbors, shown in Fig. 7, to nearly zero by dimension 5. The attractor, which is essentially five dimensional, is twisted significantly in the time delay coordinate system provided by the vectors $\mathbf{y}(n)$, and it requires seven dimensions to completely unfold its intersections with itself.

$$\bar{\lambda}_a(L)=\frac{1}{N_S}\sum_{k=1}^{N_S}\lambda_a(\mathbf{y}(k),L), \tag{9}$$

for N_S starting locations $\mathbf{y}(k)$. These quantities are shown in Fig. 11 and then in an enlarged view in Fig. 12 for type I chaos. There are clearly two positive Lyapunov exponents, one zero exponent, which is characteristic of the dynamics of differential equations [2], and four negative exponents. Their sum is negative, as it should be, and for this sample of type I chaos that sum is approximately -1.28 in units of inverse τ_s. The behavior of the $\lambda_a(L)$ as seen here is quite typical of average local Lyapunov exponents as seen in model dynamical systems [2]. The largest exponent starts high compared to

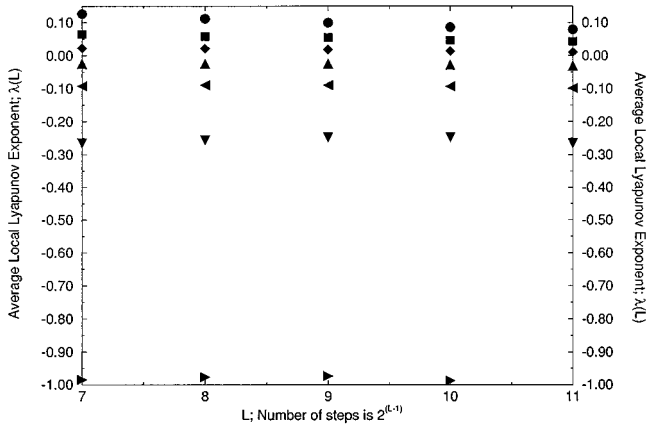

FIG. 12. A blowup of the average local Lyapunov exponents shown in Fig. 11. The zero exponent and the two positive exponents are much clearer in this view of the data.

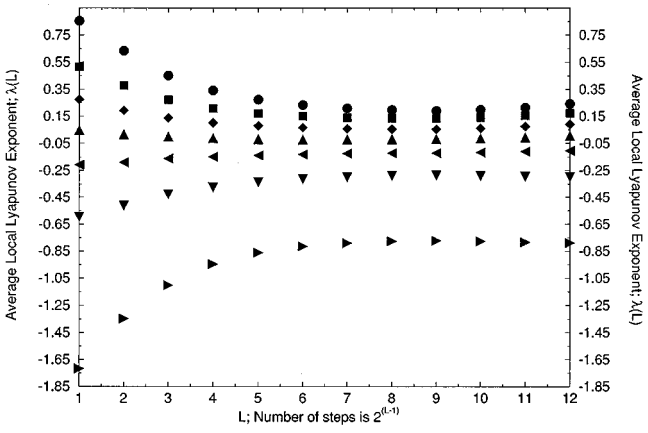

FIG. 13. The average local Lyapunov exponents for our sample of type II chaos. A global and a local dimension of $d_E = d_L = 7$ have been used, and 120 000 points from the time series were utilized. There are three positive Lyapunov exponents and one zero exponent, indicating that differential equations describe the underlying dynamics.

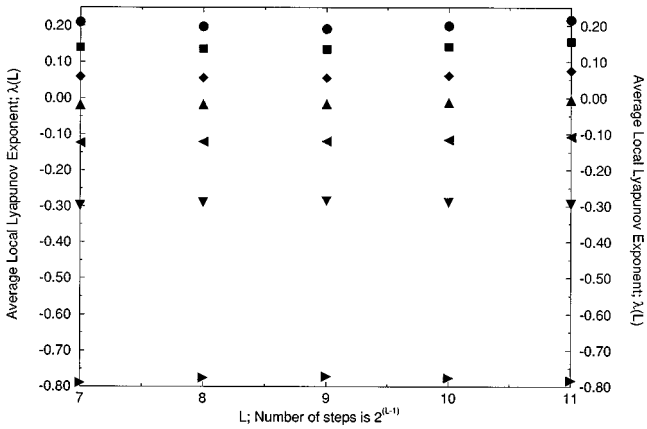

FIG. 14. A blowup of the average local Lyapunov exponents shown in Fig. 13. The zero exponent and the three positive exponents are much clearer in this view of the data.

its value for large L, while the negative exponents start low and rise. The zero exponent begins negative and rises toward zero or rides very close to zero for all L. Each of the average local exponents shown here changes sign when the eigenvalues of the Oseledec matrix is evaluated backward in time. This is as it should be for real exponents, and this supports our choice of $d_E = d_L = 7$ from local false neighbors.

The same calculations were done on our data set for type II chaos. From the outset we must recall that when a data set is noisy, the evaluation of Lyapunov exponents may well be uncertain [2]. The origin of the uncertainty is in the severely ill conditioned nature of the Oseledec matrix, which serves to amplify any numerical errors in the individual Jacobians composing it. With noise the determination of neighboring distances and the local map from which we read off the Jacobian are sure to lead to small errors in each Jacobian. This will lead to real uncertainties in the Lyapunov exponents, especially the negative exponents [2].

Nonetheless, we have evaluated the average local Lyapunov exponents for these data, and in Fig. 13 we present the results of this calculation. Figure 14 shows an enlargement of the average exponents for large numbers of steps along the trajectory after a perturbation to the known orbit. From these figures we see that the largest exponent is about twice that of the largest type I exponent. This would lead immediately to increased unpredictability, as we saw in the evaluation of local false nearest neighbors. Also we see three positive exponents, which would be connected to the apparent inability to control the type II chaos with the methods that have been tried [3,4]. It is reassuring that one of the exponents is zero, so we again have a set of differential equations describing the source of these data.

A check of our calculations is presented in Table I where we display the values of the average local Lyapunov exponents evaluated after 2048 steps along the attractor following a perturbation. The exponents forward and backward are presented for each of the types of chaos we consider. If the exponents are true, then their signs should reverse when time is reversed. Within errors commensurate with the experimental resolution, this reversal is seen to be quite accurate. We also note that the largest Lyapunov exponents for type II chaos is nearly 3 times that for type I chaos. The substantially greater unpredictability of type II chaos is quantified by this.

IV. PREDICTING IN RECONSTRUCTED PHASE SPACE

Even without knowledge of the dynamical equations for the laser system we can use the information we have acquired so far to make predictive models for the laser intensity evolution. The method [2] utilizes the compactness of the attractor in $\mathbf{y}(n)$ space by noting that we have knowledge of the evolution of whole phase-space neighborhoods into whole neighborhoods later in time. We can use this to make local models of this evolution and then use these models as interpolating rules for the evolution of new phase-space points near the attractor. We use only local polynomial models although other basis sets than polynomials are certainly quite useful. Indeed, in this work we use only local linear models as we have substantial amount of data and thus good coverage of the attractor; namely, every neighborhood is rather well populated.

The idea is that locally on the attractor we find the N_B neighbors $\mathbf{y}^{(r)}(n)$; $r = 1,2,\ldots N_B$ of each point $\mathbf{y}(n)$ and make

TABLE I. Average local Lyapunov exponents for laser chaos.

Average Lyapunov exponents at $L = 2048$ $d_E = 7$; $d_L = 7$			
Type I chaos		Type II chaos	
Forward	Backward	Forward	Backward
0.080	−0.089	0.244	−0.271
0.041	−0.043	0.172	−0.188
0.008	−0.013	0.091	−0.113
−0.033	0.029	0.0069	−0.0017
−0.102	0.096	−0.104	0.091
−0.278	0.264	−0.298	0.274
−1.017	1.020	−0.788	0.733

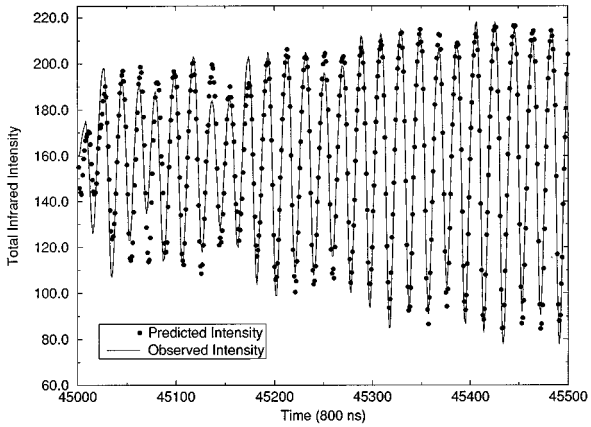

FIG. 15. Predicted and observed total infrared intensity for type I chaos. The predictions are made using local linear maps in the reconstructed phase space of the attractor. The coefficients for the local maps are learned from the first 40 000 points of the data set, and then predictions are made starting with point 45 000. All predictions are made in $d_E=d_L=7$. This figure shows the result of predicting ahead $\Delta t=10\tau_s$. The largest global Lyapunov exponent is about $\frac{1}{12}$ in units of τ_s, so we do not expect to be able to accurately predict ahead much more than this amount.

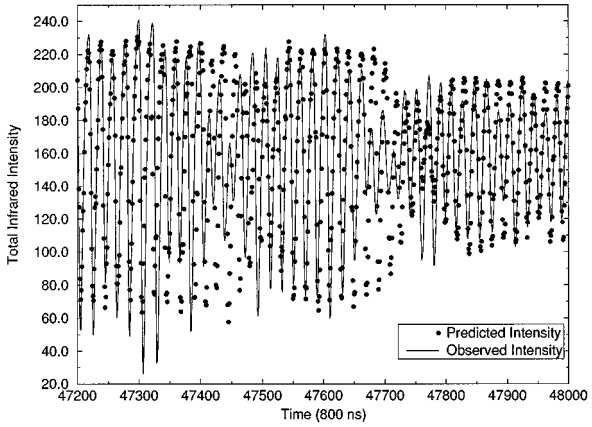

FIG. 16. Predicted and observed total infrared intensity for type I chaos. The predictions are made using local linear maps in the reconstructed phase space of the attractor. The coefficients for the local maps are learned from the first 40 000 points of the data set, and then predictions are made starting with point 45 000. All predictions are made in $d_E=d_L=7$. This figure shows the result of predicting ahead $\Delta t=50\tau_s$. The predictions are much worse than for $\Delta t=10\tau_s$ since we are trying to predict beyond τ_s/λ_1.

the local linear model $\mathbf{y}(r;n+1)=\mathbf{A}_n+\mathbf{B}_n\cdot\mathbf{y}^{(r)}(n)$ where $\mathbf{y}(r;n+1)$ is the point to which $\mathbf{y}^{(r)}(n)$ goes in one time step. The coefficients $\mathbf{A}_n$ and $\mathbf{B}_n$ are determined by minimizing at each time location n

$$\sum_{r=0}^{N_B}|\mathbf{y}(r;n+1)-\mathbf{A}_n-\mathbf{B}_n\cdot\mathbf{y}^{(r)}(n)|^2, \tag{10}$$

where $r=0$ means $\mathbf{y}(n)$ itself and $\mathbf{y}(n+1)$. When we have a new point $\mathbf{z}(k)$ on or near the attractor, we seek the nearest-neighbor $\mathbf{y}(Q)$ among all the data in the set we used to determine the $\mathbf{A}_n$ and the $\mathbf{B}_n$. The predicted point $\mathbf{z}(k+1)$ is then

$$\mathbf{z}(k+1)\approx\mathbf{A}_Q+\mathbf{B}_Q\cdot\mathbf{z}(k). \tag{11}$$

This works remarkably accurately within the limits of prediction dictated by the largest Lyapunov exponent λ_1. When we try to predict beyond the instability horizon, that is for times much greater than

$$\frac{\tau_s}{\lambda_1}, \tag{12}$$

our prediction should rapidly lose accuracy.

In Fig. 15 we show an example of this prediction technique for type I chaos based on a total data set of 60 000 points. 40 000 points were used to determine the local polynomial coefficients $\mathbf{A}_n$ and $\mathbf{B}_n$, then predictions were made 10 steps ahead from the point $n=45\,000$ to the point $n=55\,000$. The results for points $n=45\,500$ are shown in the figure. The predictions are shown as solid symbols while the observed data are shown as the solid line. In units of τ_s the largest Lyapunov exponent is approximately $\lambda_1\approx0.08$, so we should be able to make accurate predictions out to twelve or so steps beyond any starting location on the attractor. The computations reported in Fig. 15 support this quite well. In Fig. 16 we show the result of the same calculation but now predicting ahead 50 steps instead of the 10 just shown. There are clearly regions where the predictability is rather good, but also regions where the method starts to fail quite visibly. Figure 17 shows a region where the predictability remains quite good, while Fig. 18 is a region where predictability is quite reduced for this large time step ahead of a known point. These results are consistent with the wide variation of *local*

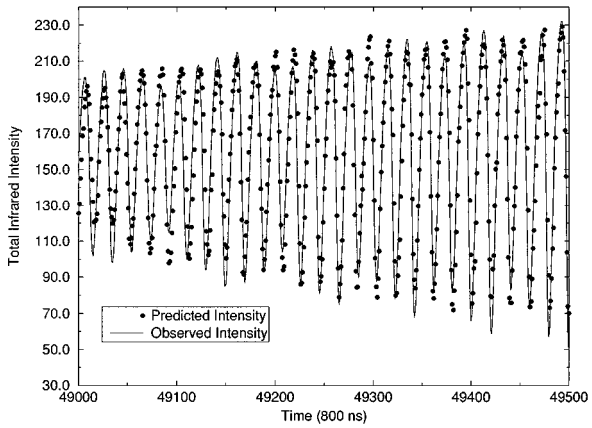

FIG. 17. Predicted and observed total infrared intensity for type I chaos. The predictions are made using local linear maps in the reconstructed phase space of the attractor. The coefficients for the local maps are learned from the first 40 000 points of the data set, and then predictions are made starting with point 45 000. All predictions are made in $d_E=d_L=7$. This figure shows the result of predicting ahead $\Delta t=50\tau_s$. This is a region of phase space where the predictions are much better than would be expected from the values of the *global* Lyapunov exponents.

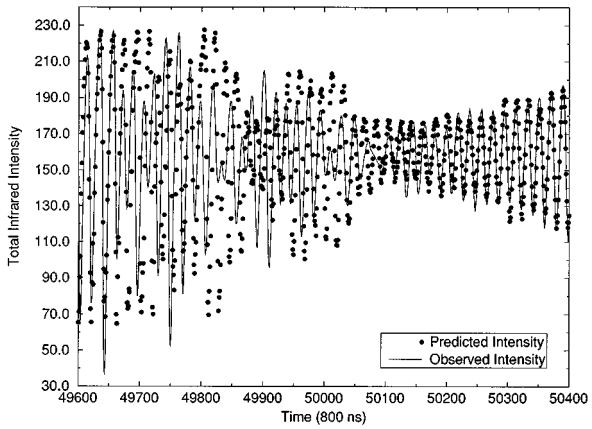

FIG. 18. Predicted and observed total infrared intensity for type I chaos. The predictions are made using local linear maps in the reconstructed phase space of the attractor. The coefficients for the local maps are learned from the first 40 000 points of the data set, and then predictions are made starting with point 45 000. All predictions are made in $d_E = d_L = 7$. This figure shows the result of predicting ahead $\Delta t = 50\tau_s$. This is a region of phase space where the predictions are much worse than when we tried to predict ahead only $\Delta t = 10\tau_s$.

Lyapunov exponents on an attractor [2]. Pushing this even further we show in Fig. 19 the result of trying to predict ahead 100 steps beyond a new point near the attractor. Clearly predictability has been reduced substantially, as it should be reduced in a chaotic system.

We have also used this method for moving about the attractor for type II chaotic data. In Fig. 20 we display the result of learning on 40 000 points of this data set and then predicting ahead from point 45 000. The predictions ten steps ahead compared to the observed values for $I(t)$ are shown in this figure for time steps 55 000–57 500. In Fig. 21 we have

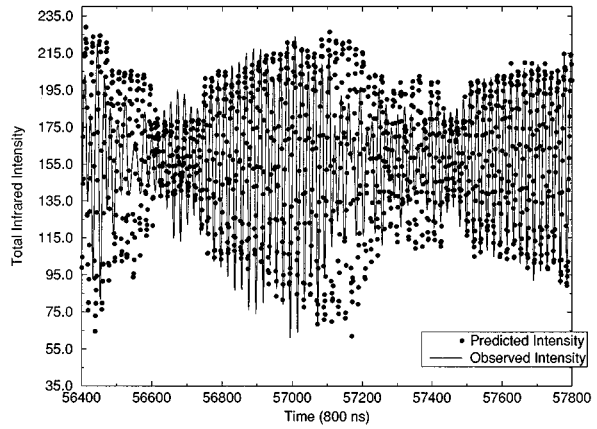

FIG. 19. Predicted and observed total infrared intensity for type I chaos. The predictions are made using local linear maps in the reconstructed phase space of the attractor. The coefficients for the local maps are learned from the first 40 000 points of the data set, and then predictions are made starting with point 45 000. All predictions are made in $d_E = d_L = 7$. This figure shows the result of predicting ahead $\Delta t = 100\tau_s$. The predictions are much worse than above since we are trying to predict quite a bit beyond τ_s/λ_1.

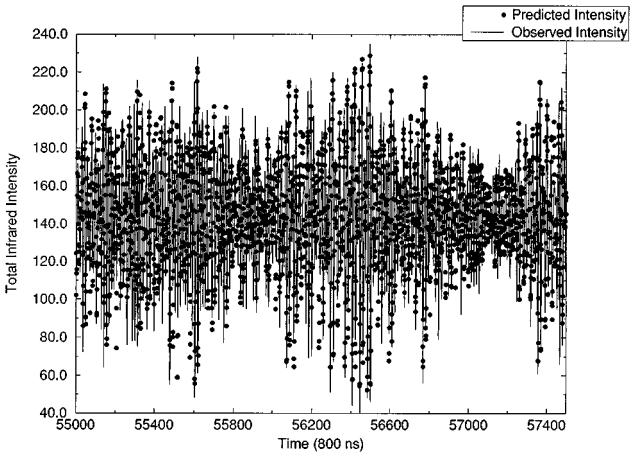

FIG. 20. Predicted and observed total infrared intensity for type II chaos. The predictions are made using local linear maps in the reconstructed phase space of the attractor. The coefficients for the local maps are learned from the first 40 000 points of the data set, and then predictions are made starting with point 45 000. All predictions are made in $d_E = d_L = 7$. This figure shows the result of predicting ahead $\Delta t = 10\tau_s$.

enlarged the region between steps 55 000 and 55 500 so one may see that the quality of these predictions is not as good as we saw in working with type I data. This is as it should be. Figure 22 is another enlargement of the type II predictions making much the same point about the quality of the predictions. Finally in Fig. 23 we have predictions 50 time steps ahead for type II chaos. Again 40 000 points were used for learning the local linear maps used for prediction. The quality of the predictions here has become quite poor.

V. A MODEL FOR THE PROCESS

The salient features of our data and data analysis can be summarized as follows: (i) There are only a few modes of infrared light in the laser cavity. Indeed, the number in this experiment was determined to be three by observing the infrared light in a Fabry-Pérot interferometer. (ii) These infrared modes with wavelength $\lambda \approx 1.064\ \mu$m couple through the KTP crystal to green light at wavelength $\lambda \approx 0.532\ \mu$m. The

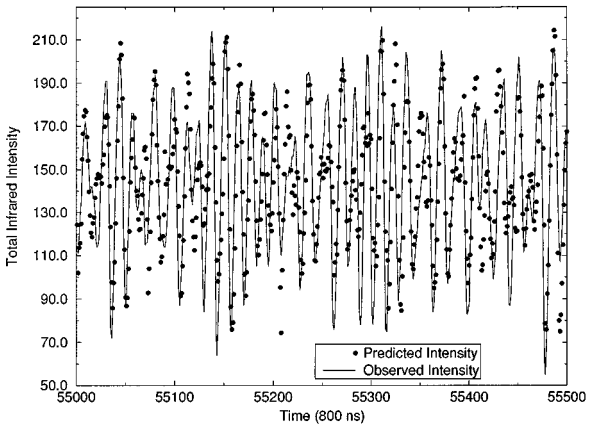

FIG. 21. An enlargement of the predictions for $\Delta t = 10\tau_s$ in type II chaos.

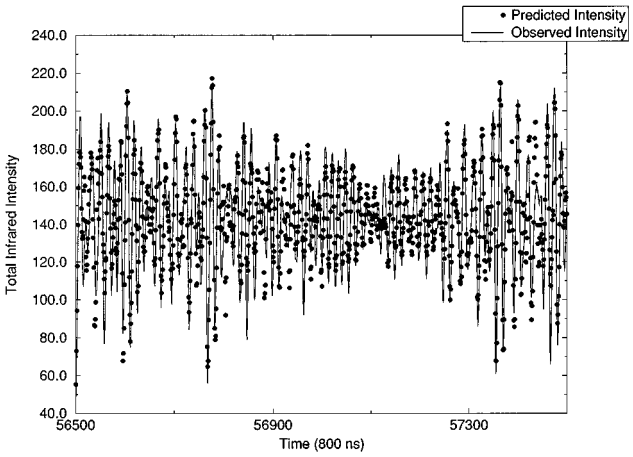

FIG. 22. An enlargement of the predictions for $\Delta t = 10\tau_s$ in type II chaos.

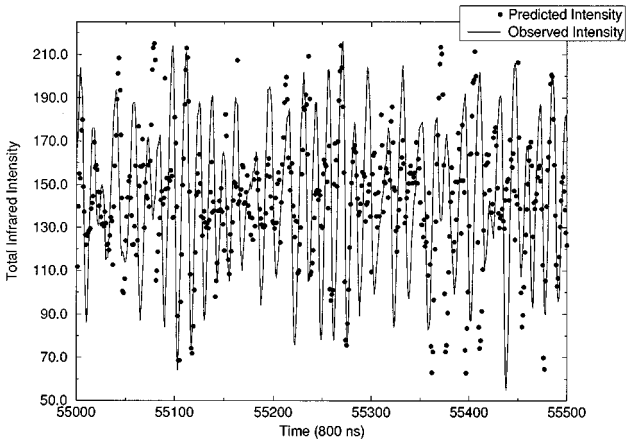

FIG. 23. Predicted and observed total infrared intensity for type II chaos. The predictions are made using local linear maps in the reconstructed phase space of the attractor. The coefficients for the local maps are learned from the first 40 000 points of the data set, and then predictions are made starting with point 45 000. All predictions are made in $d_E = d_L = 7$. This figure shows the result of predicting ahead $\Delta t = 50\tau_s$.

amount of green light depends on the polarization of the infrared modes. If all modes are polarized parallel to each other, the production is small. If one mode is polarized perpendicular to the other two, the production of green light is strongly enhanced. (iii) When green light is produced, it exits the cavity through one of the mirrors, which is transmitting at that wavelength. The same mirror reflects the infrared light, so green is not a cavity mode. (iv) The infrared production in the active medium is pumped by a diode laser at $\lambda \approx 810$ nm.

The model we suggest neglects the detailed dynamics of the active medium, treating it as a continuum of two-level quantum systems distributed over the length of the laser cavity. Each two-level system has an upper level $|u(z)\rangle$ and a lower level $|l(z)\rangle$ with energy difference $\hbar\omega_A$. The dynamical variables for the atomic levels are the usual Pauli spin operators $S_3(z,t)$ and $S_\pm(z,t)$. They satisfy

$$[S_3(z), S_\pm(z')] = \pm 2 S_\pm(z)\,\delta(z-z') \tag{13}$$

and

$$[S_+, S_-(z')] = S_3(z)\,\delta(z-z'), \tag{14}$$

at equal times. We have annihilation and creation operators $a_m^\dagger(t)$ and $a_m(t)$ for the M infrared modes with the usual equal-time Bose commutation relations

$$[a_m^\dagger, a_n] = \delta_{mn}. \tag{15}$$

The indices m or n refer both to the mode and to the polarization. Green light is treated as a scalar field with creation and annihilation operators $g^\dagger(t)$ and $g(t)$ satisfying

$$[g^\dagger, g] = 1, \tag{16}$$

at equal times. Of course, the green light has polarization and probably more than a single mode, but these were not measured in our experiment, so we will not address those properties of the green light. Further, we shall see that the green modes are damped out so strongly relative to the infrared that they trail or are fully determined by the infrared dynamics. The green modes play a "nondynamical" role in these experiments.

The loss of infrared light in mode i will be treated in a conventional manner as a coupling to a "reservoir" of radiation modes $c_{ik}^\dagger$ and c_{ik}, which are also bosons. The green loss is modeled in the same way using radiation modes $c_{gk}^\dagger$ and c_{gk} with a higher coupling constant indicating a higher loss rate. The loss mechanisms for the atomic levels, essentially the electromagnetic modes responsible for spontaneous emission from the upper level $|u(z)\rangle$, are labeled $b_k(z)$ and $b_k^\dagger(z)$. These modes of the electromagnetic field serve solely to describe the loss mechanism as seen in the lasing system.

The Hamiltonian for the system is written as

$$\begin{aligned} H = & \sum_{i=1}^{M} \hbar\omega_i a_i^\dagger a_i + \hbar\omega_g g^\dagger g + i\hbar \sum_{i,j=1}^{M} [\kappa_{ij} a_i^\dagger a_j^\dagger g - \kappa_{ij}^* g^\dagger a_j a_i] + \sum_{i=1}^{M} \sum_k [\hbar\Omega_{ik} c_{ik}^\dagger c_{ik} + \gamma_{ik}^* a_i c_{ik}^\dagger + \gamma_{ik} c_{ik} a_i^\dagger] \\ & + \sum_k [\hbar\Omega_{gk} c_{gk}^\dagger c_{gk} + \gamma_{gk}^* g c_{gk}^\dagger + \gamma_{gk} c_{gk} g^\dagger] + \int_0^L \left[\frac{\hbar\Omega_A}{2} S_3(z) + i\hbar \sum_{i=1}^{M} [\sigma_i S_+(z) a_i \sin k_i z - \sigma_i^* a_i^\dagger S_-(z) \sin k_i z] \right. \\ & \left. + i\hbar \sum_k [\Gamma_k(z) S_+(z) b_k(z) - \Gamma_j^*(z) b_k^\dagger(z) S_-(z) + \hbar\omega_k^{\mathrm{sp}}(z) b_k^\dagger(z) b_k(z)] \right] dz. \end{aligned} \tag{17}$$

In this expression, the ω_i are the frequencies associated with the infrared modes; ω_g is the green frequency: $\omega_g \approx 2\omega_i$. Ω_{ik} and Ω_{gk} are the frequencies of the reservoir modes coupled to the infrared mode i and green mode, respectively, and $\omega_j^{\rm sp}(z)$ are the reservoir frequencies for the atomic loss mechanism at location z. The nonlinear coupling κ_{ij} comes from the second-order susceptibility of the KTP crystal. It is dependent on the polarizations of the infrared photons. The γ_{ik} tell us the strength of the coupling of the infrared modes to their reservoirs, the γ_{gk} is the coupling of the green mode to its reservoir, and $\Gamma_k(z)$ is the coupling of the atomic levels to their reservoirs at location z. There is no pumping of the two-level system included in this model Hamiltonian. It is added to the Heisenberg equation of motion for $S_3(z)$ as a constant rate of population inversion.

From the Hamiltonian we determine the Heisenberg equations of motion for the various operators and perform standard reservoir approximations [10] to get a model that can be integrated numerically. Using this model we have successfully reproduced type I chaos and this will be discussed in a future paper [14].

VI. SUMMARY AND CONCLUSIONS

We have determined the properties of an Nd:YAG laser with an intracavity KTP crystal when the laser is in an operating mode producing chaotic fluctuations of the total output infrared intensity. While this chaotic intensity has been studied in the past [1] we have brought to the analysis several tools not previously applied to this problem: (i) average mutual information, (ii) global and local false nearest neighbors, and (iii) determination of the full spectrum of local and global Lyapunov exponents.

Using these new tools we have been able to start from the observation of a scalar—total infrared light intensity $I(t)$—and deduce the number of active degrees of freedom in the laser cavity, evaluate the spectrum of Lyapunov exponents, and make ''black box'' predictive models that use no physics associated with the laser but, nonetheless, allow accurate prediction on or near the observed attractor. We have also suggested a physical quantum mechanical model for which numerical studies are still being performed. We anticipate that the resulting equations contain all of the features, including the nonlinear characteristics, of our data.

In laser operation our system has exhibited two general classes of dynamical activity. Each has three infrared modes lasing in the cavity. For the chaotic operating state we have labeled type I chaos all three of the infrared modes are polarized parallel to each other. In this state the production of green light is small—less than 1 μW in our cavity—and the quantum noise, as revealed by global false nearest neighbors, is small. This means that the macroscopic dynamical system involving infrared intensity and perhaps gain is well described by a few differential equations with some small associated noise terms. In the operating state we have called type II chaos there are also three modes of infrared light though two are parallel polarized and one is polarized perpendicular to the other two. The production of green light is much stronger—about 25 μW—and the influence of noise, as revealed by global false nearest neighbors is much larger than for type I chaos.

Our emphases in this paper have been twofold: First, we have been concerned with the application of a general program of nonlinear time series analysis [2] to the specific problem of the dynamics of a certain solid state laser system. The description of that set of tools as they are used in a practical setting can be no better given than using our work here as an example. Second, we have been interested in the properties of this specific laser system and in determining characteristics that describe its behavior. In that regard the local and global Lyapunov exponents are important aspects of the dynamics of this laser source. In fact, by establishing that type II chaos has both high false neighbors at high embedding dimension and correspondingly a higher value for the largest positive global Lyapunov exponent, we have found that type II chaos is quantitatively less predictable and, consistent with earlier experience [3,4] less controllable than type I chaos. This provides a consistent picture for these operating conditions for this laser system.

The dynamics of the laser has been described by coupled macroscopic differential equations for the intensities $I_k(t)$ of the individual infrared longitudinal modes and their corresponding gains $G_k(t)$ [1,11]

$$\tau_c \frac{dI_k}{dt} = \left[G_k - \alpha - \epsilon g I_k - 2\epsilon \sum_{j \neq k} \mu_{jk} I_j \right] I_k , \tag{18}$$

$$\tau_f \frac{dG_k}{dt} = \gamma - \left[1 + I_k + \beta \sum_{j \neq k} I_j \right] G_k . \tag{19}$$

$\tau_c \approx 0.2$ ns is the cavity round trip time, and $\tau_f \approx 240$ μs is the fluorescence lifetime of the Nd^{3+} ions. The cavity loss $\alpha \approx 0.01$. γ is the small signal gain related to the pump rate. β is the cross-saturation parameter which we take to be the same for all mode pairs and is somewhat less than unity for this laser system. The efficiency of the nonlinear doubling process is defined by ϵ. For the KTP crystal $\epsilon \approx 10^{-5}$. g is a geometric factor $0 \leqslant g \leqslant 1$, which depends on the relative orientation between the fast axes of the strongly birefringent KTP and the very weakly birefringent Nd:YAG; it also depends on the lengths of the crystals. The factor μ_{jk} accounts for the change in geometry when the modes j and k have varying polarizations. If the modes have the same polarization $\mu_{jk} = g$, otherwise $\mu_{jk} = 1 - g$. This factor determines the relative amount of green light produced by second harmonic versus sum frequency generation for different polarization configurations of the laser modes. These equations are macroscopic descriptions of the laser that do not account for any noise sources.

The description of the laser system using these ''classical'' equations has been quite successful. The next task we will pursue is to reconcile the quantum-mechanical equations derived from the model above with these classical equations. Clearly the operator equations contain more information about the system: they deal with infrared mode amplitudes rather than intensities, so phases are included, and they contain both fluctuations associated with the quantum mechanics of the operators and associated with the loss of green light from the cavity.

In part these effects will enlarge our understanding of the sum frequency generation of green light in its semiclassical description [12] as well as numerous interesting issues associated with the classical-quantum correspondence [13] as exhibited in this particular system.

All these are important effects, and our future work on this system will now deal with numerical approximations to the Heisenberg equations which are consistent with the observed data as analyzed by our tools as well as establishing in a quantitative way the role of the fluctuations we have identified. Good values for all the relevant parameters such as the green-infrared couplings κ_{ij} and the damping rates will be part of the tasks, and we shall return to this in our next paper [14].

ACKNOWLEDGMENTS

Z.G. and R.R. thank B. Kennedy, K. Wiesenfeld, and R. Fox for many helpful discussions on the material discussed here. R. R. acknowledges support from NSF Grant No. ECS-9114232. C.L. and H.D.I.A. thank the members of INLS for numerous discussions on this subject. Their work was supported in part by the U.S. Department of Energy, Office of Basic Energy Sciences, Division of Engineering and Geosciences, under Contract No. DE-FG03-90ER14138, and in part by the Army Research Office (Contract No. DAAL03-91-C-052), and by the Office of Naval Research (Contract No. N00014-91-C-0125), under subcontract to the Lockheed/Sanders Corporation.

[1] C. Bracikowski and R. Roy, Chaos **1**, 49 (1991).
[2] H. D. I. Abarbanel, R. Brown, J. J. Sidorowich, and Lev Sh. Tsimring, Rev. Mod. Phys. **65**, 1331 (1993).
[3] R. Roy, T. W. Murphy, Jr., T. D. Maier, Z. Gills, and E. R. Hunt, Phys. Rev. Lett. **68**, 1259 (1992).
[4] Z. Gills, C. Iwata, R. Roy, I. B. Schwartz, and I. Triandaf, Phys. Rev. Lett. **69**, 3169 (1992).
[5] A. V. Oppenheim and R. W. Schafer, *Digital Signal Processing* (Prentice-Hall, Englewood Cliffs, NJ, 1989), Chap. 5.
[6] A. M. Fraser and H. L. Swinney, Phys. Rev. A **33**, 1134 (1986); A. M. Fraser, IEEE Trans. Info. Theory **35**, 245 (1989).
[7] Matthew B. Kennel, R. Brown, and H. D. I. Abarbanel, Phys. Rev. A **45**, 3403 (1992).
[8] H. D. I. Abarbanel and M. B. Kennel, Phys. Rev. E **47**, 3057 (1993).
[9] V. I. Oseledec, Trudy Mosk. Mat. Obsc. **19**, 197 (1968) [Moscow Math. Soc. **19** 197 (1968)].
[10] H. J. Kimble, in *Fundamental Systems in Quantum Optics*, edited by J. Dalibard, J.-M. Raimond, and J. Zinn-Justin, Les Houches, Session LIII (North-Holland, Amsterdam, 1992), pp. 545–674.
[11] T. Baer, J. Opt. Soc. Am. B **3**, 1175 (1986).
[12] P. D. Drummond, K. J. McNeil, and D. F. Walls, Opt. Acta **28**, 211 (1981).
[13] R. F. Fox and T. C. Elston, Chaos **3**, 313 (1993).
[14] C. Liu, H. D. I. Abarbanel, Z. Gills, and R. Roy (unpublished).