ABSTRACT

Title of Dissertation: LOW DIMENSIONAL CHAOS: PHASE SYNCHRONIZATION AND INDETERMINATE BIFURCATIONS

Romulus Breban, Doctor of Philosophy, 2003

Dissertation directed by: Professor Edward Ott
Department of Physics

We address two problems of both theoretical and practical importance in dynamical systems: Phase Synchronization of Chaos in the Presence of Two Competing Periodic Signals, and Saddle-Node Bifurcations on Fractal Basin Boundaries.

LOW DIMENSIONAL CHAOS:
PHASE SYNCHRONIZATION AND
INDETERMINATE BIFURCATIONS

by

Romulus Breban

Dissertation submitted to the Faculty of the Graduate School of the
University of Maryland, College Park in partial fulfillment
of the requirements for the degree of
Doctor of Philosophy
2003

Advisory Committee:

Professor Edward Ott, Chairman/Advisor
Professor Rajarshi Roy
Professor Thomas M. Antonsen
Professor James A. Yorke
Professor Brian R. Hunt

DEDICATION

To my parents

ACKNOWLEDGEMENTS

I am grateful to Prof. Edward Ott for constant guidance through my Ph.D. research, and for all lessons of life that came along with it. I am also grateful to Prof. Helena E. Nusse for guidance and collaboration, especially for opening my eyes to the beautiful and delicate world of fractal basin boundaries.

During my Ph.D. years I was part of the 'Maryland Chaos Group', and I thank all my fellow students Doug Armstead, Jon-Won Kim, Juan Restrepo, Mike Oczkowski, Seung Jong Baek, Xing Zheng and Su Li for many interesting discussions about dynamical systems and physics, and especially for their friendship.

I address special thanks to my wife Mihaela who stood by me throughout these years.

TABLE OF CONTENTS

LIST OF TABLES

LIST OF FIGURES

Chapter 1

Introduction

The field of Dynamical Systems and Chaos has experienced a highly nonlinear growth over the past few decades. Besides major progress in topics as bifurcations, crises, basin boundaries and others, laying the mathematical foundations, new fields such as Control and Synchronization of Chaos have emerged, addressing more practical questions.

This Ph.D. thesis contains a little bit of both. The second Chapter is devoted to a problem in Phase Synchronization of Chaos. Phase Synchronization of Chaos occurs as a result of a weak interaction between dynamical systems. From the practical point of view, it offers a new valuable method of testing interdependence between time series, which found numerous applications in Neuroscience and Communications. The second Chapter discussed the situation where two periodic systems compete to entrain a chaotic oscillator. Phase Synchronization of Chaos is a common phenomenon that is expected to play a major role in disentangling the secrets of nature.

The third Chapter presents a more theoretical problem with far reaching practical implications, which is the Saddle-Node Bifurcation on a Fractal Basin Boundary. At first, it may seem puzzling that a saddle-node bifurcation may occur exactly on a fractal basin boundary, which is a zero Lebesgue measure set. However, not only does this type of bifurcation happen, but it also relates to the Wada property of fractal basins, and it is a quite common occurrence in dynamical systems. From the practical point of view, this problem belongs to the study of what happens when an attracting periodic orbit is lost due to a parameter variation through a saddle-node bifurcation. When the saddle-node bifurcation occurs on a fractal basin boundary, the attracting periodic orbit collides with an unstable periodic orbit embedded in its own basin boundary (i.e., the basin boundary of the attracting periodic orbit). The fate of an orbit following the location of the pre-bifurcation periodic attractor, as the system parameter drifts, is indeterminate. It becomes difficult (or impossible) to predict what is the final state reached by the orbit drifting past the saddle-node bifurcation. The study presented here characterizes this difficulty analytically, and opens the road for experiments detecting saddle-node bifurcations on fractal basin boundaries.

Chapter 2

Phase Synchronization of Chaos in the Presence of Two Competing Signals

2.1 Preliminaries

Phase synchronization of chaos has attracted much attention due to its applicability to a wide range of situations including laser, plasma, fluid and biological experiments. Synchronization of chaotic attractors with the phase of a periodic externally coupled signal has been studied theoretically [1, 2, 3, 4, 5] and demonstrated experimentally [6, 7]. Phase synchronization of coupled chaotic systems has also been studied [8, 9, 10, 11, 12, 13].

In order to define phase synchronism, assume that we are given two signals a and b where both possess an oscillatory character, such that phases $\phi_a(t)$ and $\phi_b(t)$ can, by some appropriate means, be defined for the two signals. Here the phases $\phi_{a,b}(t)$ are assumed to be continuous in time (i.e., they are not taken modulo 2π), so that, if, for two times $t_2 > t_1$, we have $\phi_{a,b}(t_2) - \phi_{a,b}(t_1) = 2N\pi$, then we say that the phase $\phi_{a,b}$ has executed N counter-clockwise rotations between time t_1 and time t_2. (Thus, $\phi_{a,b}$ is defined on the real line rather than on $[0, 2\pi]$. This is refered to as the "lift" of the angle.)

Two types of phase synchronism can be distinguished: *strong phase synchronism* and *weak phase synchronism*. In terms of the difference $\triangle\phi(t) = \phi_a(t) - \phi_b(t)$, there is strong phase synchronism between the signals a and b if

$$-K \leq \triangle\phi(t) - \phi_0 \leq K$$

for some constants K and ϕ_0 (typically $K \sim \pi$) and all time t. Thus, $| \triangle \phi|$ does not increase without bound. In weak phase synchronism $| \triangle \phi|$ may become arbitrarily large with increasing time, but the behavior of $\triangle\phi(t)$ as a function of time manifests correlations between the two phases (examples will be given subsequently).

In this chapter we consider the case where two periodic signals compete to entrain a chaotic oscillator. There are several possible motivations for this study. First, there may be real situations where a chaotic dynamical system simultaneously receives inputs from two distinct periodic systems (e.g., a neuron receiving signals from two other neurons). Second, the study of a signal with two frequencies can be regarded as a next step from the single frequency case in obtaining an understanding of phase synchronization of chaos by signals with nontrivial frequency power spectra (Sec. 4). Third, this situation is a generalization of the problem in which two periodic signals compete to entrain a nonlinear *periodic* oscillator.

2.2 Model Dynamical System

We consider a specific model system consisting of a *modified* chaotic Roessler [14] oscillator coupled to a two frequency input signal, $s(t)$. If we denote the regular Roessler system by $d\mathbf{x}/dt = \mathbf{R}(\mathbf{x})$, where $\mathbf{x}^T = (x(t), y(t), z(t))$, then our modified (undriven) system is [4] $d\mathbf{x}/dt = f(\mathbf{x})\mathbf{R}(\mathbf{x})$, where f is a scalar function of $\mathbf{x}$ that is positive in the region of the chaotic attractor. This modification of the Roessler system does not change the topology of the trajectory curves followed by orbits in phase space, but it does modify the speed with which orbits move along these curves. The motivation for doing this [4] is that the original Roessler system displays a frequency spectrum with a near-delta-function-like feature, corresponding to the average period for an orbit to circulate arround the attractor. This type of behavior is typically not present or expected in the experimental studies [6, 7, 9, 10, 11, 12, 13]. By our modification, we introduce enhanced dispersion in the time for an orbit to circulate around the attractor, and hence the width in the Fourier peak. We take $f(\mathbf{x}) = 1 + \sigma(r^2 - \bar{r}^2)$, $\sigma = 0.002$,

$r^2 = x^2 + y^2$, with $\bar{r}$ equal to the time average of r for the unmodified and unentrained Roessler system ($\bar{r} = 5.037$) [15]. Our model system becomes [4]:

$$\begin{aligned} dx/dt &= -[1 + 0.002(r^2 - \bar{r}^2)](y + z), \\ dy/dt &= [1 + 0.002(r^2 - \bar{r}^2)](x + 0.25y) + s(t), \\ dz/dt &= [1 + 0.002(r^2 - \bar{r}^2)][0.90 + z(x - 6.0)], \end{aligned} \tag{2.1}$$

where

$$s(t) = A_1 \cos(\omega_1 t) + A_2 \cos(\omega_2 t), \tag{2.2}$$

and we have chosen the parameters of the Roessler system so that it is in the so-called phase coherent regime (i.e., the x-y projection of the trajectory of the chaotic system with $A_1 = A_2 = 0$ continually circles around $x = y = 0$, and the x-y projection of the attractor appears to be shaped like an annulus with $x = y = 0$ in the hole of the annulus). Our main goals in this study are to examine the illustrative system (2.1), (2.2) in different regimes, and to delineate and explain the various types of observed phenomena. We conjecture that the phenomena we observe for the system (2.1), (2.2) are typical for general oscillatory chaotic systems subject to two frequency external driving.

From studies of the phase synchronism of chaos by a single sinusoidal signal, $s_0(t) = A_0 \sin \omega t$, $\omega = 2\pi/T$, [3, 4] it is known that the parameter space given by the amplitude A_0 and period T of the signal typically displays a tongue-shaped region where the phase of the attractor locks with the phase of the periodic signal (i.e., perfect phase synchronism), as shown schematically in Fig. 2.1(a). For the purpose of the subsequent discussion we also note that the two frequency entraining signal (2.2) can

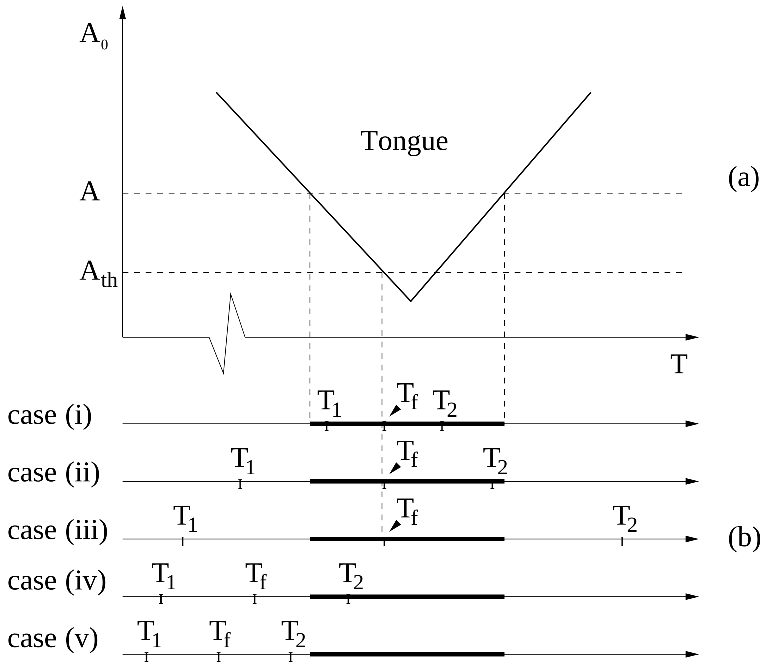

Figure 2.1: (a) Schematic of the parameter space A_0-T for the case where there is a single sinusoidal signal, $s(t) = A_0 \cos(2\pi t/T)$, coupled to the Roessler system. (b) Illustration of various cases for the situation in which a signal, consisting of the sum of two equal amplitude sinusoids, $s(t) = A\cos(2\pi t/T_1) + A\cos(2\pi t/T_2)$, is coupled with the Roessler system [$T_1 < T_2$, $T_f = 2T_1T_2/(T_1 + T_2)$]. The bold horizontal lines represent the range of T over which phase synchronism occurs for a single sinusoidal signal of amplitude $A_0 = A$.

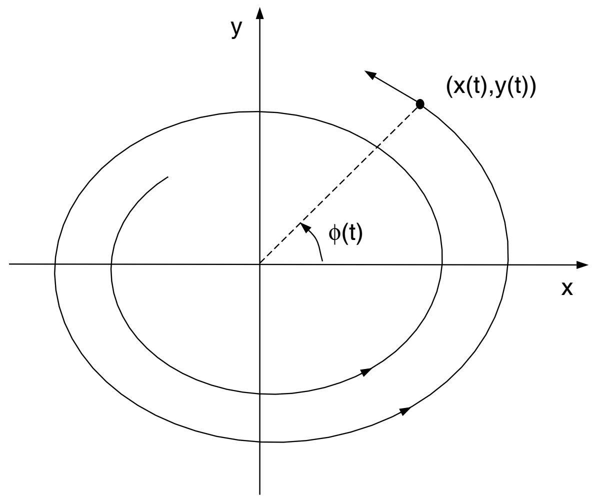

Figure 2.2: Graphical illustration of the definition of geometrical phase $\phi(t)$ for a chaotic orbit.

be written in an alternate form,

$$s(t) = (A_1 + A_2)\cos[(\omega_1 + \omega_2)t/2]\cos[(\omega_1 - \omega_2)t/2]$$
$$+(A_2 - A_1)\sin[(\omega_1 + \omega_2)t/2]\sin[(\omega_1 - \omega_2)t/2]. \tag{2.3}$$

In most of our numerical work we have considered the case of equal amplitudes $A_1 = A_2 = A = 0.06$. (Later we will discuss the case where A_1 and A_2 are different.) From (2.3), the entraining signal $s(t)$ can be regarded as a modulated wave, a "fast wave" at the mean frequency

$$\omega_f = (\omega_1 + \omega_2)/2$$

modulated by a "slow wave" at the frequency

$$\omega_s = (\omega_1 - \omega_2)/2,$$

where, for $A_1 = A_2 = A$, the modulating slow wave is

Table 2.1: Parameter values T_1 and T_2

case	T_1	T_2
i)	5.95	5.99
ii)	5.90	5.99
iii)	5.00	7.40
iv)	5.00	5.99
v)	5.00	5.50

$$\hat{A}(t) = 2A\cos[(\omega_1 - \omega_2)t/2].$$

In our numerical experiments $(\omega_1 + \omega_2) \gg (\omega_1 - \omega_2) > 0$. Three periods will prove relevant: $T_{1,2} = 2\pi/\omega_{1,2}$ and $T_f = 2\pi/\omega_f = 2T_1T_2/(T_1 + T_2)$. The geometrical phase of an orbit (Fig. 2.2) is given by $\tan\phi(t) = [y(t)/x(t)]$ where the relevant branch of $\tan\phi(t) = [y(t)/x(t)]$ is determined by the previously mentioned definition of $\phi(t)$ as continuous in t; see Sec. 1. We investigate how $\phi(t)$ is related to the phases of the sinusoidal signals $\phi_{1,2} = \omega_{1,2}t$ as well as to the phase based on the mean frequency $\phi_f = \omega_f t$. As in previous studies, the phase differences,

$$\triangle\phi_{1,2,f}(t) = \phi(t) - \phi_{1,2,f},$$

are used to test phase synchronism between the chaotic orbits of our driven Roessler system (2.1), (2.2) and one of the three phases ϕ_1, ϕ_2 or ϕ_f.

We note that synchronism at $\phi_f = \frac{1}{2}(\omega_1+\omega_2)t$ can be viewed as a special case of the general situation where $l\phi$ synchronizes with $m\phi_1+n\phi_2$, where l, m and n are integers. In this framework, synchronism with ϕ_f corresponds to $l = 2$ and $m = n = 1$.

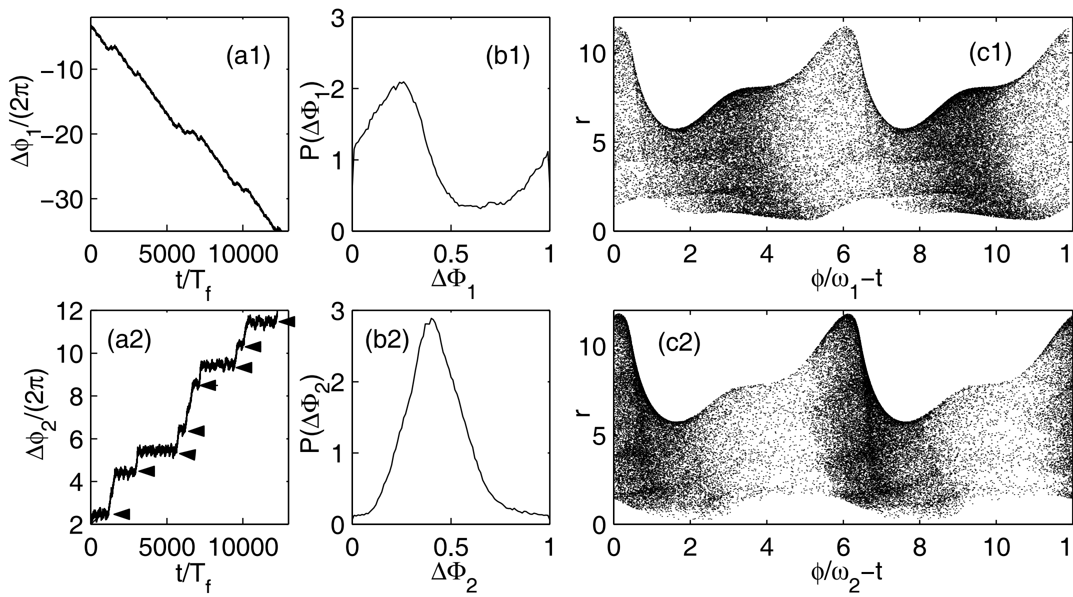

Figure 2.3: (a1,2) Difference between the geometrical phase of the attractor ϕ and the phase of the first/second sinusoidal signal $\phi_{1,2}$ versus time t/T_f. (b1,2) Histogram approximations of the distribution functions $P(\triangle\Phi_{1,2})$, where $\triangle\Phi_{1,2} = [\triangle\phi_{1,2}/(2\pi)]$ modulo 1. (c1,2) Stroboscopic sections at times $t = nT_{1,2}$ (n is an integer) through the perturbed Roessler attractor, Eqs. (1) and (2).

2.3 Results

We now report and discuss results of computations for several different choices of the parameters T_1 and T_2. These results serve to illustrate the main qualitative behaviors that we have found. In particular, we consider the five sets of parameter values given in Table 1. For each of the parameter sets of Table 1 the disposition of the values T_1, T_2 and T_f with respect to the tongue of perfect phase synchronism for a single frequency driving signal is illustrated schematically in Figs. 2.1(b)-2.1(f). We first give a detailed account for case (i) followed by brief descriptions of the results for the other cases.

2.3.1 Case (i)

In this case there are clear intervals of time, lasting many rotations of ϕ or $\phi_{1,2,f}$ [note that $(\omega_1 - \omega_2)/\omega_f \ll 1$], when ϕ is entrained by ϕ_2. In such a time interval, the fluctuation of $\triangle\phi_2/(2\pi)$ is limited to within a narrow range, while

$$\triangle\phi_1/(2\pi) = \triangle\phi_2/(2\pi) - (\omega_1 - \omega_2)t/(2\pi)$$

decreases with time at an average rate $(\omega_1 - \omega_2)/(2\pi)$. This behavior is seen in Figs. 2.3(a1) and 2.3(a2) which show $\triangle\phi_1/(2\pi)$ and $\triangle\phi_2/(2\pi)$ versus $\phi_f/(2\pi) = t/T_f$ over a range representing over 10^4 rotations of ϕ_f. Refering to Fig. 2.3(a2), plateaus representing locking of ϕ to ϕ_2 are clearly evident and are indicated on the figure by arrowheads (the longest of these plateaus represents approximately 500 rotations of ϕ_f). We also note that each plateau is centered at a value of $\triangle\phi_2/(2\pi)$ that is larger than that for the previous plateau by an integer. That is, ϕ slips relative to ϕ_2 by an integer number of complete rotations between plateaus. [By the arrowheads in Fig. 2.3(a2) we have considered a plateau to exist if it is at least as wide as $T_s/2 = 2\pi/(\omega_1 - \omega_2)$, i.e., half the period of the slow wave.] Refering to Fig. 2.3(a1), we see that the graph of $\triangle\phi_1/(2\pi)$ versus $\phi_f/(2\pi) = t/T_f$ appears to consist of intervals of approximate linear decrease (with superposed fluctuations) at a slope $-(\omega_1 - \omega_2)/\omega_f$ separated by glitches. The intervals of time corresponding to apparent linear decrease of $\triangle\phi_1/(2\pi)$ coincide with the plateaus of $\triangle\phi_2/(2\pi)$, while the glitches in $\triangle\phi_1/(2\pi)$ coincide with the time intervals between the plateaus of $\triangle\phi_2/(2\pi)$. Alternatively, one may consider these glitches to be narrow plateaus of $\triangle\phi_1/(2\pi)$. A close examination of Fig. 2.3(a1) also shows that the average values of $\triangle\phi_1/(2\pi)$ corresponding to these narrow plateaus differ by integers. ϕ slips

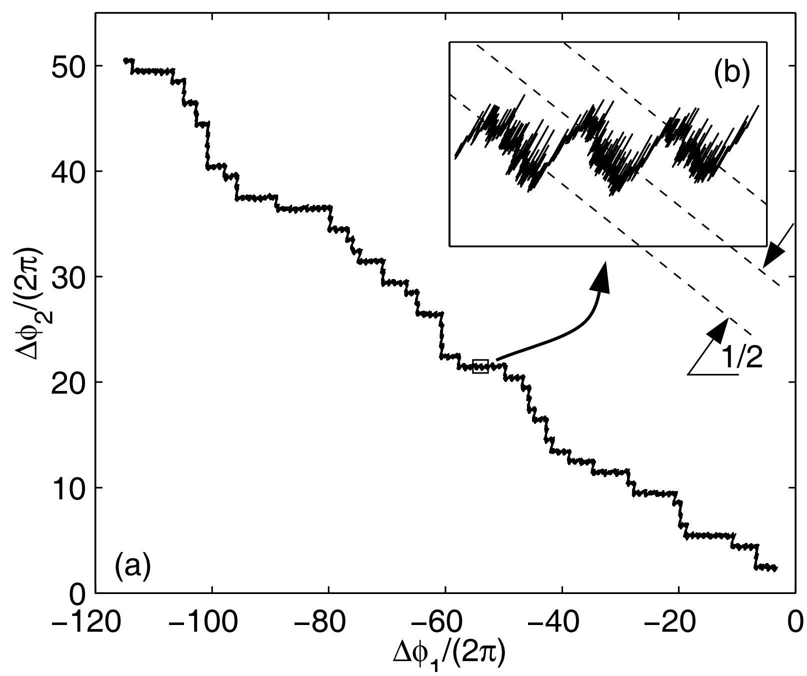

Figure 2.4: (a) $\triangle\phi_2/(2\pi)$ versus $\triangle\phi_1/(2\pi)$. The staircase-like structure indicates that there are alternating time intervals in which ϕ is locked to either ϕ_1 or to ϕ_2. (b) Detail of Figure 4(a).

relative to ϕ_1 by an integer number of complete rotations between plateaus. Thus in the competition between ϕ_1 and ϕ_2 to entrain ϕ, there are intervals when ϕ_1 wins and intervals when ϕ_2 wins, but, overall, ϕ_2 is a stronger entrainer that ϕ_1. This is also indicated in Figs. 2.3(a1) and 2.3(a2) by the fact that, in the same time interval, $\triangle\phi_1$ goes through more that 30 rotations, while $\triangle\phi_2$ only goes through 9. [The relative entraining strengths of ϕ_1 and ϕ_2 depend on the locations of T_1 and T_2 within the tongue in Fig. 2.1(a).] Figure 2.4(a) graphs $\triangle\phi_1$ versus $\triangle\phi_2$. The staircase-like structure shows that when $\triangle\phi_1$ varies, $\triangle\phi_2$ is aproximately constant and vice versa; the approximately horizontal portions of the graph correspond to plateaus of $\triangle\phi_2$ and the approximately vertical portions correspond to plateaus of $\triangle\phi_1$. This supports the picture whereby we can think of the chaotic oscillator as making transitions between two states of locking with the phases $\phi_{1,2}$ of the competing signals.

Figures 2.3(b1) and 2.3(b2) show histogram approximations of the probability distributions of $\triangle\Phi_1 \equiv \triangle\phi_1/(2\pi)$ modulo 1 and, respectively, $\triangle\Phi_2 \equiv \triangle\phi_2/(2\pi)$ modulo 1 [16]. The purpose of these figures is to demostrate that statistically significant correlations between ϕ and $\phi_{1,2}$ can be found. That is, each of the phases ϕ_1 and ϕ_2 weakly synchronize the chaotic attractor. [In the absence of any coupling between ϕ and $\phi_{1,2}$ these graphs would be flat, $P(\triangle\Phi_{1,2}) = 1$.]

Figures 2.3(c1) and 2.3(c2) show stroboscopic surfaces of section at the periods T_1 and, respectively, T_2. For each point on a long trajectory we plot r versus $[\phi \,\text{modulo}\, 4\pi]/\omega_{1,2} - t$. This gives a picture of the density of the strobed points on the attractor. Both Figs. 2.3(c1) and 2.3(c2) show alternating regions of high and low density of points. (One should imagine an infinite periodic chain of such regions from which we only plotted two periods.) The high density regions represent regions where the orbit spends a long time. The low density regions are regions that the orbit traverses relatively fast. Therefore, the plateaus of Figs. 2.3(a1) [respectively, Fig. 2.3(a2)] correspond to regions with high density in Fig. 2.3(c1) [respectively, Fig. 2.3(c2)]. The times when ϕ slips with respect to $\phi_{1,2}$ generate regions of low density. The fact that, when Fig. 2.3(c1) has a low density region, Fig. 2.3(c2) has a high density region corresponds to the fact that when ϕ slips with respect to ϕ_1, it locks with respect to ϕ_2.

We now consider the possibility of phase synchronism of our system with the fast wave phase $\phi_f = \omega_f t$. Using (2.3), we think of $s(t)$ as a sinusoid entraining at the period T_f (the period of the fast wave) slowly modulated at the period T_s (the period of the slow wave). When the amplitude $\hat{A}$ of the fast wave becomes smaller than

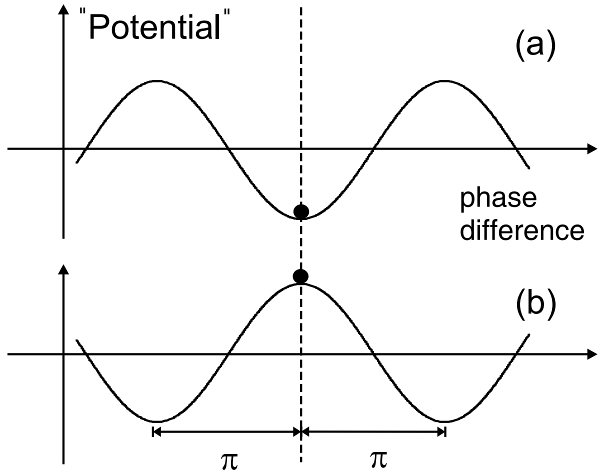

Figure 2.5: Particle in sinusoidal potential: (a) at minimum potential, (b) at maximum potential after the sign change of the potential.

the threshold A_{th} set by the synchronization tongue at T_f [see Fig. 2.1(a)], the chaotic attractor tends [17] to lose synchronization and slip with respect to the phase of the fast wave. The synchronization condition $|\hat{A}| > A_{th}$ implies that the attractor tends to lose synchronization as $\hat{A}$ drops below A_{th} but tends to synchronize as $\hat{A}$ decreases through $-A_{th}$. Let τ denote the duration of a time interval during which $|\hat{A}| < A_{th}$ in a slow wave period T_s. If we consider the phase $\phi'(t)$ of the free running Roessler system (i.e., Eqs. (2.1)-(2.3) with $\hat{A} = 0$), then, during the time τ, the phase difference $\phi'(t) - \omega_f t$ is found to change by less than π. Thus, during a time interval τ, we expect that there is not sufficient time for $\triangle\phi_f$ to drift as much as 2π before resynchronizing after $|\hat{A}|$ exceeds A_{th}. Thus, we anticipate that slips of $\triangle\phi_f$ are solely due to the change in sign of $\hat{A}$. These slips are expected to be $\pm\pi$. In order to see this, we make a crude analogy, and consider a particle in the vicinity of a potential minimum in a sinusoidal potential [analogous to the fact that the phase $\phi(t)$ is in the vicinity of $\phi_f(t)$]; see Fig. 2.5(a). If we now change the sign of the potential, then the particle finds itself at the top of a

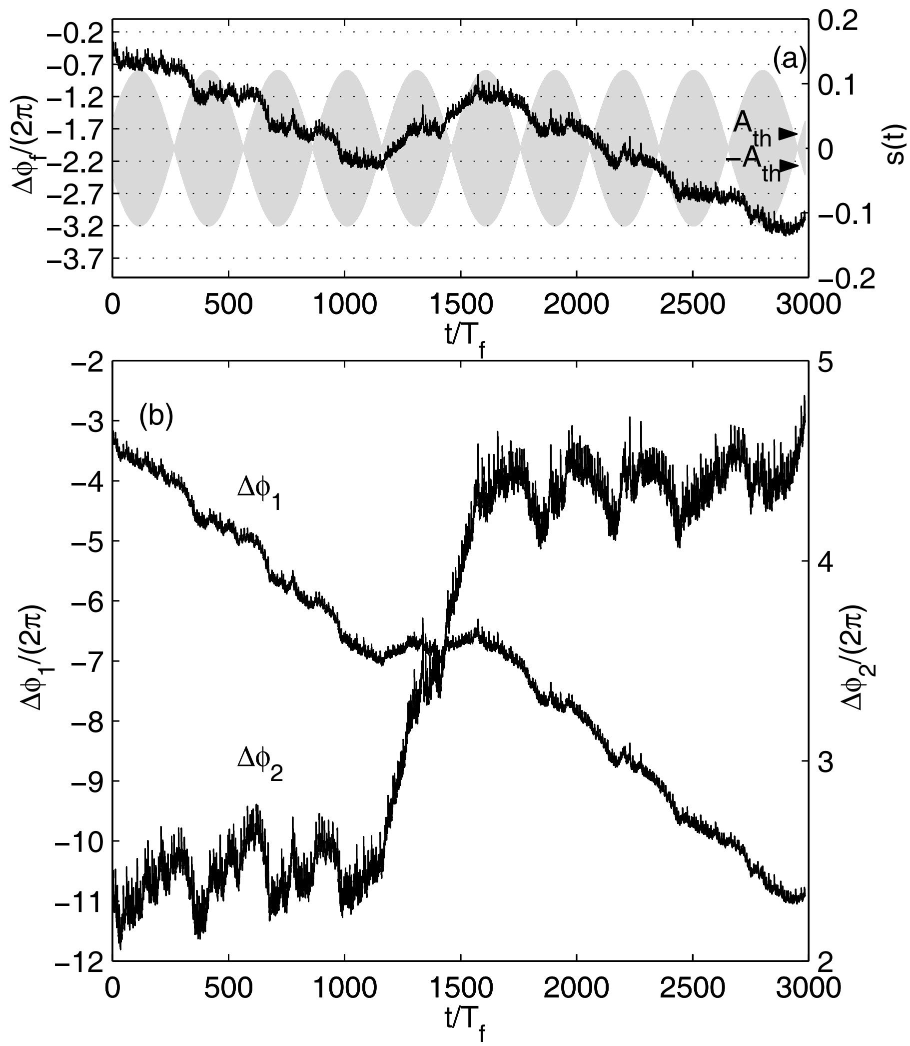

Figure 2.6: (a) Detail of how $\triangle\phi_f$ switches between slipping down to slipping up with the entraining modulating slow wave indicated by the grey background. (b) Detail of how the chaotic attractor switches between locking to ϕ_2 and locking to ϕ_1. The time axes in (a) and (b) coincide.

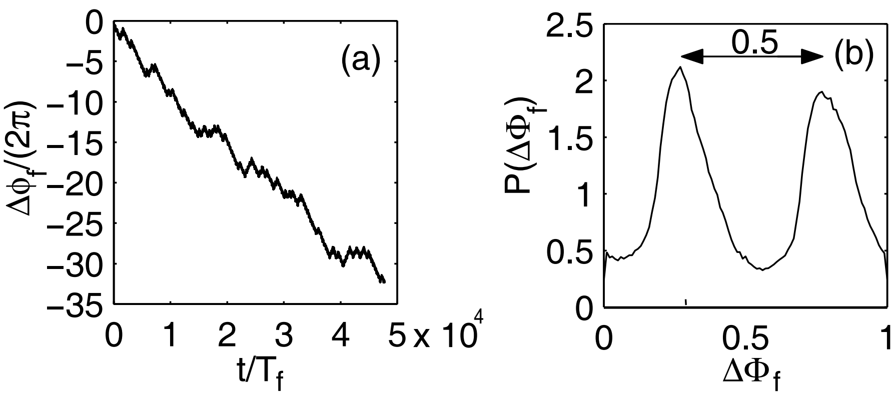

Figure 2.7: (a) $\triangle\phi_f/(2\pi)$ versus time t/T_f. (b) Histogram approximation of the distribution function $P(\triangle\Phi_f)$, where $\triangle\Phi_f = [\triangle\phi_f/(2\pi)]$, modulo 1.

potential hill, and (assuming appropriate friction) will take some time to evolve to one of the adjacent minima situated at a phase of the potential that is $\pm\pi$ away [Fig. 2.5(b)]. By these considerations, we can expect that the graph of $\triangle\phi_f$ versus time will display plateaus of synchronization and slips of π up or down occurring twice every period of the slow wave. This is illustrated in Fig. 2.6(a) which shows how $\triangle\phi_f$ varies with time for several periods of the slow wave. $s(t)$ is plotted as the grey background for convenience. To guide the eye, dotted horizontal lines separated by a change of π in $\triangle\phi_f$ are drawn through the plateaus.

Fig. 2.6(b) displays $\triangle\phi_1(t)$ and $\triangle\phi_2(t)$ in the same range of time as in Fig. 2.6(a). Comparison of Figs. 2.6(a) and 2.6(b) reveals that time intervals of locking with the phase of the fast wave ϕ_f with π slips down correspond with the time of locking with the phase ϕ_2, while time intervals of locking with the phase of the fast wave ϕ_f with π slips up correspond with the time of locking with the phase ϕ_1. We also remark that when $\triangle\phi_f$ has a plateau, $\triangle\phi_{1,2}$ drifts slowly at the rate $\omega_f - \omega_{1,2}$ with superimposed fluctuations. During the time when $\triangle\phi_f$ slips down, $\triangle\phi_2$ may stay locked. For example, see Fig. 2.6(b) which shows $\triangle\phi_2/(2\pi)$ to be in a plateau for $t/T_f < 1100$. In this range, the graph of $\triangle\phi_2/(2\pi)$ versus t/T_f has a roughly sawtooth-like structure, with an upward drift with slope $(\omega_f - \omega_2)/\omega_f$ during the plateaus of $\triangle\phi_f$ and rapid decrease between the plateaus of $\triangle\phi_f$.

Fig. 2.7(a) shows $\triangle\phi_f$ over a much longer time scale than is plotted in Fig. 2.6(a). Refering to Fig. 2.4(a) and noting that $[\triangle\phi_1/(2\pi) + \triangle\phi_2/(2\pi)]/2 = \triangle\phi_f/(2\pi)$ and $[\triangle\phi_1/(2\pi) - \triangle\phi_2/(2\pi)]/2 = t/T_s$, it is seen that a $\pi/4$ rotation and a change of scale converts Fig. 2.4(a) to Fig. 2.7(a). In these coordinates [Fig. 2.4(a)], the jumps along the horizontal and vertical axis are integers. A close inspection of Fig. 2.4(a) reveals that the plateaus of $\triangle\phi_2/(2\pi)$ plotted versus $\triangle\phi_1/(2\pi)$ are not entirely flat. They have a rough saw-tooth structure in which saw-tooth segments of slope -1 correspond to the times of locking of ϕ with ϕ_f (such locking implies $\triangle\phi_1 + \triangle\phi_2 \sim$ constant). This is indicated by the blow-up, Fig. 2.4(b), where dashed lines of slope -1 going through the plateaus of locking with ϕ_f are shown. These lines are separated by $1/2$, corresponding to the $\pm\pi$ slips in Fig. 2.6(a). Fig. 2.7(b) shows a histogram approximation of the probability distribution of $\triangle\Phi_f \equiv \triangle\phi_f/(2\pi)$ modulo 1 demonstrating that the phase of the attractor ϕ weakly synchronizes with ϕ_f. The probability distribution of $\triangle\Phi_f$ in Fig. 2.7(b) has two maxima 0.5 apart because $\triangle\phi_f$ undergoes $\pm\pi$ jumps. This is in contrast with the probability distributions for $\triangle\phi_{1,2}$ which have only one maximum, corresponding to the fact that $\triangle\phi_{1,2}$ undergo $\mp 2\pi$ jumps, respectively.

2.3.2 Other Cases

Case (ii) In this case, (A_1, T_1) is outside the single sinusoid synchronization tongue, while (A_2, T_2) and $(\hat{A}, T_f)$ are inside. Histogram approximations to the distributions

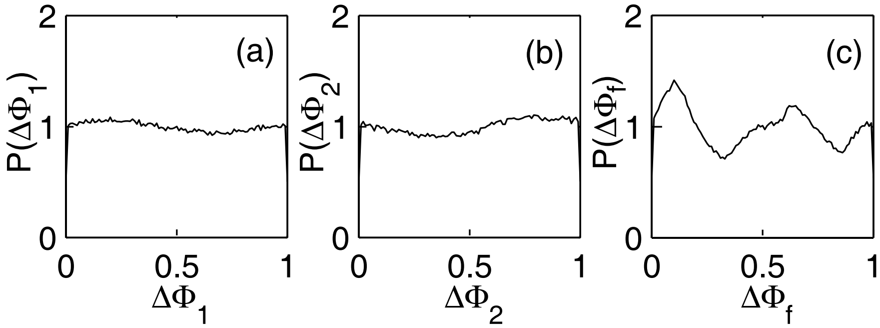

Figure 2.8: Results for case (iii). (a,b,c) Histogram approximations of the distribution functions $P(\triangle\Phi_{1,2,f})$, where $\triangle\Phi_{1,2,f} = [\triangle\phi_{1,2,f}/(2\pi)]$ modulo 1.

$P(\triangle\Phi_1)$, $P(\triangle\Phi_2)$, and $P(\triangle\Phi_f)$ (figures not included) all differ significantly from the flat distribution and look very similar with those for case (i) in Figs. 2.3(b1), 2.3(b2), and 2.7(b), respectively. Thus, some degree of synchronization of the chaotic system with all phases ϕ_1, ϕ_2, and ϕ_f is manifest. In addition, plots of $\triangle\phi_1$, $\triangle\phi_2$, and $\triangle\phi_f$ versus time (not included) look very similar to those in Figs. 2.3(a1), 2.3(a2), and 2.7(a). However, in comparison with case (i), there is significantly enhanced tendency for synchronization with phase ϕ_2 as opposed to ϕ_1. The plateaus of $\triangle\phi_2$ are longer (in average) and the plateaus of $\triangle\phi_1$ are shorter than in the case (i). $\triangle\phi_f/(2\pi)$ still shows plateaus of synchronization but mostly slips up corresponding with the fact that almost all the time ϕ_2 synchronizes the orbit. At times, $\triangle\phi_f/(2\pi)$ also shows slips down, corresponding to the little bit of time the orbit spends synchronized with ϕ_1.

Case (iii) In this case, $(A_1 = A_2 = A, T_f)$ lies inside the single sinusoid synchronization tongue, while (A_1, T_1) and $(\hat{A}, T_f)$ are outside. Figure 2.8 shows histogram approximations to the distributions $P(\triangle\Phi_1)$ [Fig. 2.8(a)], $P(\triangle\Phi_2)$ [Fig. 2.8(b)], and $P(\triangle\Phi_f)$ [Fig. 2.8(c)]. We see that $P(\triangle\Phi_1)$ and $P(\triangle\Phi_2)$ are nearly flat, indicating

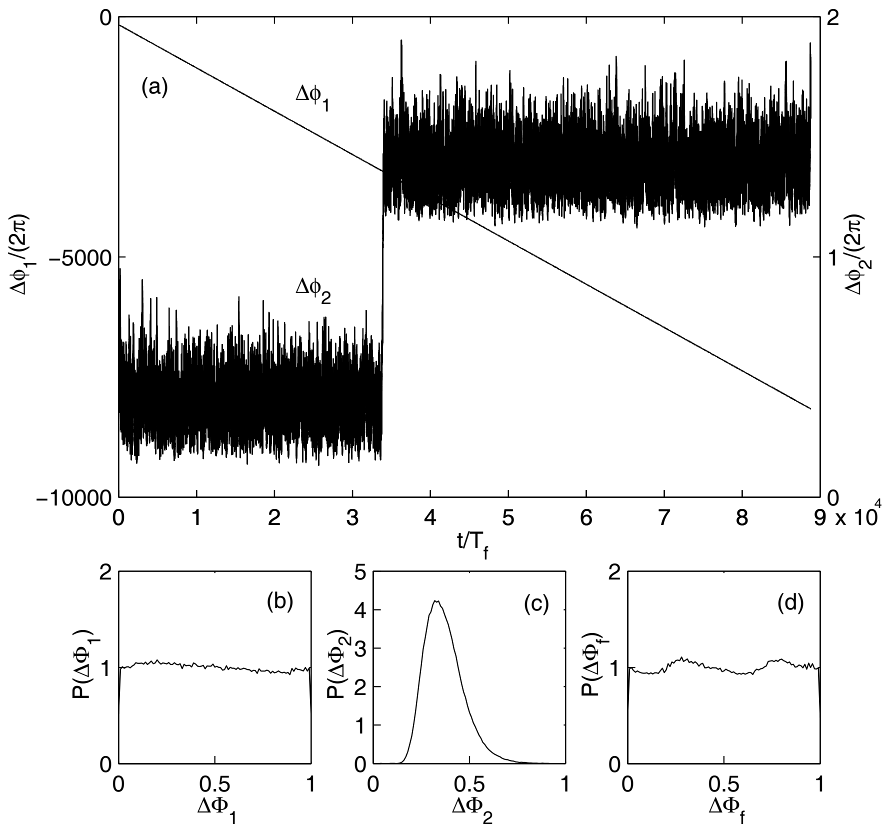

Figure 2.9: Results for case (iv). (a) $\triangle\phi_1/(2\pi)$ and $\triangle\phi_2/(2\pi)$ versus time t/T_f. (b,c,d) Histogram approximations of the distribution functions $P(\triangle\Phi_{1,2,f})$, where $\triangle\Phi_{1,2,f} = [\triangle\phi_{1,2,f}/(2\pi)]\,\mathrm{modulo}\,1$.

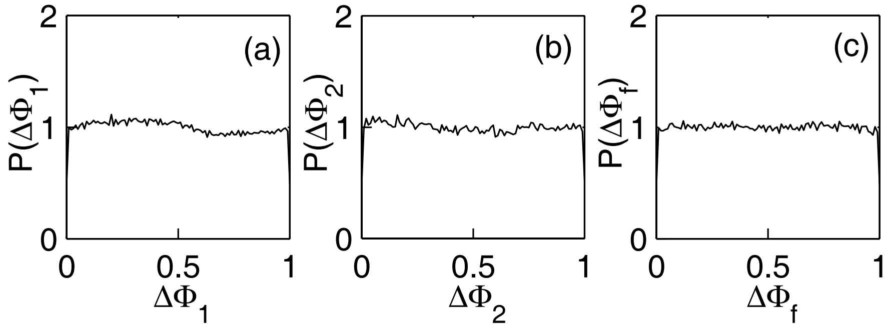

Figure 2.10: Results for case (v). (a,b,c) Histogram approximations of the distribution functions $P(\triangle\Phi_{1,2,f})$, where $\triangle\Phi_{1,2,f} = [\triangle\phi_{1,2,f}/(2\pi)]$ modulo 1.

very small, or negligible synchronization with phases ϕ_1 and ϕ_2. In contrast, $P(\triangle\Phi_f)$ shows two significant peaks separated by 0.5 in $\triangle\Phi_f$. This is similar to the plot of $P(\triangle\Phi_f)$ for case (i) shown in Fig. 2.7(b). In addition, plots of $\triangle\phi_1$ and $\triangle\phi_2$ versus time (not included) show nearly steady linear drift, while a plot of $\triangle\phi_f$ versus time (also not included) evidences periods of locking similar to Fig. 2.7(a) for case (i). Thus, for case (iii), we conclude that there is negligible synchronization of the system with the phases ϕ_1 and ϕ_2, but that there is significant synchronization with ϕ_f.

Case (iv) This case has only (A_2, T_2) inside the synchronization tongue, while (A_1, T_1) and $(\hat{A}, T_f)$ are outside. Figures 2.9(b), 2.9(c), and 2.9(a), respectively, show histogram approximations to the distributions $P(\triangle\Phi_1)$, $P(\triangle\Phi_2)$, and $P(\triangle\Phi_f)$. We remark that $P(\triangle\Phi_1)$ and $P(\triangle\Phi_f)$ are almost flat, indicating little synchronization of the chaotic system with phases ϕ_1 and ϕ_f. On the other hand, $P(\triangle\Phi_2)$ shows a big peak, suggesting synchronization with phase ϕ_2. Accordingly, the graphs of $\triangle\phi_1$ [Fig. 2.9(a)] and $\triangle\phi_f$ versus time (not included) show nearly steady linear drift, while the graph of $\triangle\phi_2$ [Fig. 2.9(a)] versus time shows very long plateaus of sychronization, indicative of strong phase synchronism (see Sec. 1). These results can be understood by noting that, by construction, case (iv) has (A_2, T_2) inside the single sinusoid synchronization tongue, while (A_1, T_1) and $(\hat{A}, T_f)$ are outside.

Case (v) In this case we have all (A_1, T_1), (A_2, T_2) and $(\hat{A}, T_f)$ outside the single sinusoid synchronization tongue. Figure 2.10 shows that histogram approximations to the distributions $P(\triangle\Phi_1)$ [Fig. 2.10(a)], $P(\triangle\Phi_2)$ [Fig. 2.10(b)], and $P(\triangle\Phi_f)$ [Fig. 2.10(c)] are all nearly flat, indicating negligible synchronization with phases ϕ_1, ϕ_2 and ϕ_f, respectively. Plots of $\triangle\phi_1$, $\triangle\phi_2$, and $\triangle\phi_f$ versus time (not included) shown nearly steady linear drift. These results are not surprising since, in this case, (A_1, T_1), (A_2, T_2), and $(\hat{A}, T_f)$ are far outside the single sinusoid synchronization tongue.

We have also investigated a few cases where A_1 and A_2 are unequal. For example, for the values of T_1, T_2 and $A_2 = 0.06$ used in case (i), we did computations for $A_1 = 0.01$ and $A_2 = 0.03$. In the former case, (A_1, T_1) is not in the synchronization tongue, and the phenomena observed are very similar to that in the case (iv) above. In the case $A_1 = 0.03$ (here (A_1, T_1) is inside the synchronization tongue) we see results similar to that in case (i), but with much reduced tendency for locking with phase ϕ_1.

2.4 Further Discussions and Conclusions

Even though our two frequency signal $s(t)$ is much simpler than entraining signals typically encountered in experiments [9, 12], we believe that it offers an important lesson regarding the understanding of synchronization by entrainers with complicated continuous frequency spectra. Data analysis of numerical and experimental results [8, 9, 12] shows that one can assign a phase to a signal (for the pupose of detecting phase synchronization of chaotic systems) by either bandpass filtering or by the use of the Hilbert transform. It has been found in experiments that the detection of phase synchronism can be enhanced by bandpass filtering [9, 12]. If we were to apply a bandpass filter to our two frequency signal $s(t)$, then, assuming a filter bandwidth less than $(\omega_1 - \omega_2)$, we would pick either the sinusoid at ω_1 or the sinusoid at ω_2, depending on the center frequency of the bandpass filter. Thus the phase of the filtered signal would be either ϕ_1 or ϕ_2. Alternatively, consider the case where we do no filtering and use the Hilbert transform technique, as advocated in [8], to produce $\tilde{s}(t)$, the complex "analytic signal" corresponding to $s(t)$. This yields

$$\tilde{s}(t) = A_1 \exp(i\omega_1 t) + A_2 \exp(i\omega_2 t).$$

The associated "Hilbert phase", ϕ_H, is:

$$\tan \phi_H = \frac{\mathrm{Im}[\tilde{s}(t)]}{\mathrm{Re}[\tilde{s}(t)]} = \frac{A_1 \sin(\omega_1 t) + A_2 \sin(\omega_2 t)}{A_1 \cos(\omega_1 t) + A_2 \cos(\omega_2 t)}.$$

For $A_1 = A_2$, this gives $\tan \phi_H = \tan[(\omega_1 + \omega_2)t/2]$, or $(\phi_H \text{ modulo } \pi) = (\phi_f \text{ modulo } \pi)$ [18]. Thus by filtering we obtain ϕ_1 or ϕ_2, while by not filtering and using the Hilbert phase we obtain ϕ_f modulo π (for $A_1 = A_2$). Which procedure is best? The answer to this question depends on circumstances. For example, in our cases (i), (ii) and (iv) synchronism with $\phi_2(t)$ is strong and clearly manifest; if a continuous spectrum had such a case, filtering might be thought to clean up the phase and make phase synchronism more apparent (as indeed has been found in some experiments [9, 12]). If, however, the situation is more like that of case (iii), where the only detectable synchronism is with ϕ_f, then narrow bandpass filtering (which yields ϕ_1 or ϕ_2) would not reveal any synchronism, while applying the Hilbert transform to the unfiltered signal would reveal synchronism.

In conclusion, we have investigated the situation in which two sinusoidal signals compete to phase synchronize a chaotic oscillator. We find and illustrate several possible outcomes of this situation:

1. Phase synchronism can be descerned to be present to some degree for both sinusoides as well as for the mean phase of the sinusoides, ϕ_f [cases (i) and (ii)].

2. Phase synchronism can be descernable only for the mean phase [case(iii)].

3. Phase synchronism is descernable only for one of the sinusoids [case(iv)].

Chapter 3

Saddle-Node Bifurcations on a Fractal Basin Boundary

3.1 Preliminaries

It is common for dynamical systems to have two or more coexisting attractors. In predicting the long-term behavior of a such a system, it is important to determine sets of initial conditions of orbits that approach each attractor (i.e., the basins of attraction). The boundaries of such sets are often fractal ([19], Chapter 5 of [20], and references therein). The fine-scale fractal structure of such a boundary implies increased sensitivity to errors in the initial conditions: Even a considerable decrease in the uncertainty of initial conditions may yield only a relatively small decrease in the probability of making an error in determining in which basin such an initial condition belongs [19, 20]. For discussion of fractal basin boundaries in experiments, see Chapter 14 of [21].

Thompson and Soliman [22] showed that another source of uncertainty induced by fractal basin boundaries may arise in situations in which there is slow (adiabatic) variation of the system. For example, consider a fixed point attractor of a map (a node). As a system parameter varies slowly, an orbit initially placed on the node attractor moves with time, closely following the location of the solution for the fixed point in the absence of the temporal parameter variation. As the parameter varies, the node attractor may suffer a saddle-node bifurcation. For definiteness, say that the node attractor exists for values of the parameter μ in the range $\mu < \mu_*$, and that the saddle-node bifurcation of the node occurs at $\mu = \mu_*$. Now assume that, for a parameter interval $[\mu_L, \mu_R]$ with $\mu_L < \mu_* < \mu_R$, in addition to the node, there are also two other attractors A and B, and that the common boundary of the basin of attractor A, attractor B and the node is a fractal basin boundary. We are interested in the typical case where, before the bifurcation, the saddle lies on the fractal basin boundary, and thus, at the bifurcation, the merged saddle-node orbit is on the basin boundary. In such a case an arbitrarily small ball about the saddle-node at $\mu = \mu_*$ contains pieces of the basins of both A and B. Thus, as μ slowly increases through μ_*, it is unclear whether the orbit following the node will go to A or to B after the node attractor is destroyed by the bifurcation. In practice, noise or round-off error may lead the orbit to go to one attractor or the other, and the result can often depend very sensitively on the specific value of the slow rate at which the system parameter varies.

We note that the study of orbits swept through an indeterminate saddle-node bifurcation belongs to the theory of dynamical bifurcations. Many authors have analyzed orbits swept through other bifurcations, like the period doubling bifurcation [23], the pitchfork bifurcation [24, 25], and the transcritical bifurcation [25]. In all these studies of the bifurcations listed above, the local structure before *and* after the bifurcation includes stable invariant manifolds varying smoothly with the bifurcation parameter (i.e., a stable fixed point that exists before or after the bifurcation, and whose location varies smoothly with the bifurcation parameter). This particular feature of the local bifurcation structure, not shared by the saddle-node bifurcation, allows for well-posed, locally defined, problems of dynamical bifurcations. The static saddle-node bifurcation has received much attention in theory and experiments [26, 27, 28], but so far, no dynamical bifurcation problems have been defined for the saddle-node bifurcation. In this study, we demonstrate that, in certain common situations, global structure (i.e., an invariant Cantor set or a fractal basin boundary) adds to the local properties of the saddle-node bifurcation and allows for well-posed problems of dynamical bifurcations.

Situations where a saddle-node bifurcation occurs on a fractal basin boundary have been studied in two dimensional Poincaré maps of damped forced oscillators [22, 29, 30]. Several examples of such systems are known [22, 30], and it seems that this is a common occurence in dynamical systems. In this chapter, we first focus on saddle-node bifurcations that occur for one parameter families of smooth one dimensional maps having multiple critical points (a critical point is a point at which the derivative of the map vanishes). Since one dimensional dynamics is simpler than two dimensional dynamics, indeterminate bifurcations can be more simply studied, without the distraction of extra mathematical structure. Taking advantage of this, we are able to efficiently investigate several scaling properties of these bifurcations. In particular, we investigate the scaling of (1) the fractal basin boundary of the static (i.e., unswept) system near the saddle-node bifurcation (Secs. 3.2.2 and 3.2.3), (2) the dependence of the orbit final destination on the sweeping rate (Sec. 3.2.4), (3) the dependence of the time it takes for an attractor to capture the swept orbit following the bifurcation on the sweeping rate (Sec. 3.2.5), and (4) the dependence of the final attractor capture probability on the noise level (Sec. 3.2.6). Following our one-dimensional investigations, we explain that these results apply to two dimensional systems. We show, through numerical experiments on the periodically forced Duffing oscillator, that the scalings we have found also apply to higher dimensional systems (Sec. 3.3).

For one-dimensional maps, a situation dynamically similar to that in which there is indeterminacy in which attractor captures the orbit can also occur in cases where there are two rather than three (or more) attractors (Sec. 3.4). In particular, we can have the situation where one attractor persists for all values of the parameters we consider, and the other attractor is a node which is destroyed via a saddle-node bifurcation on the basin boundary separating the basins of the two attractors. In such a situation, an orbit starting on the node, and swept through the saddle-node bifurcation, will go to the remaining attractor. It is possible to distinguish different ways that the orbit initially on the node approaches the remaining attractor. We find that the way in which this attractor is approached can be indeterminate.

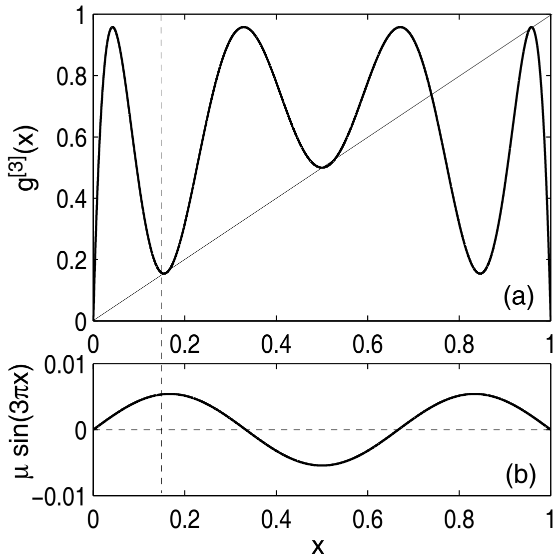

Figure 3.1: Construction of the function $f_\mu(x)$ starting with (a) the third iterate of the logistic map, $g(x) = r\,x(1-x)$, with $r = 3.832$, and adding a perturbation (b) $\mu \sin(3\pi x)$ ($\mu = 5.4 \times 10^{-3}$).

3.2 Indeterminacy in Which Attractor Is Approached

We consider the general situation of a one dimensional real map $f_\mu(x)$ depending on a parameter μ. We assume the following: (1) the map is twice differentiable with respect to x, and once differentiable with respect to μ (the derivatives are continuous); (2) f_μ has at least two attractors sharing a fractal basin boundary for parameter values in the vicinity of μ_*; and (3) an attracting fixed point x_* of the map $f_\mu(x)$ is destroyed by a saddle-node bifurcation as the parameter μ increases through a critical value μ_*, and this saddle-node bifurcation occurs on the common boundary of the basins of the two attractors.

We first recall the saddle-node bifurcation theorem (see for example [26]). If the map $f_\mu(x)$ satisfies: (a) $f_{\mu_*}(x_*) = x_*$, (b) $\frac{\partial f_{\mu_*}}{\partial x}(x_*) = 1$, (c) $\frac{\partial^2 f_{\mu_*}}{\partial^2 x}(x_*) > 0$, and (d)

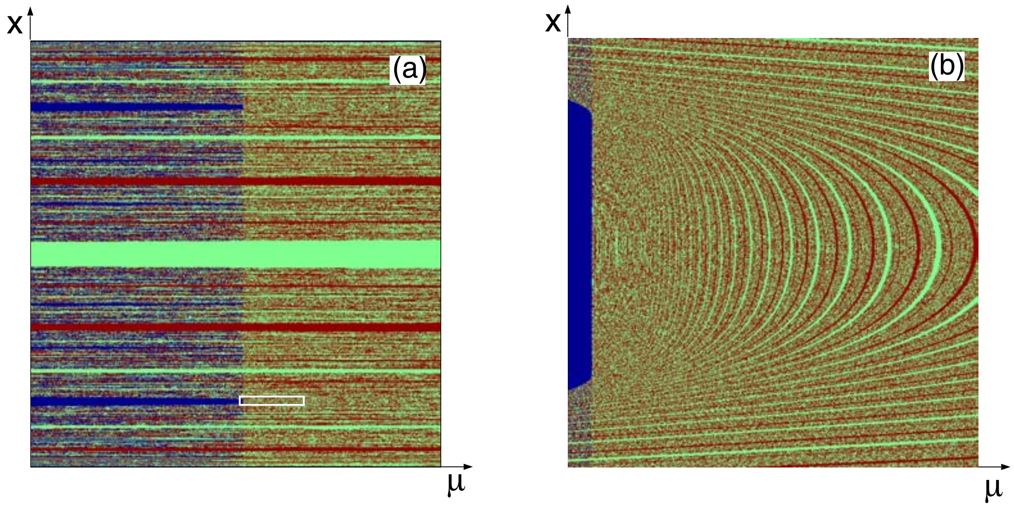

Figure 3.2: (a) Basin structure of the map f_μ versus the parameter μ on the horizontal axis ($0 \leq \mu \leq 5.4 \times 10^{-3}$ and $0 \leq x \leq 1$). The attractor having the blue basin is destroyed at $\mu \approx 2.79 \times 10^{-3}$. (b) Detail of the region shown as the white rectangle in Fig. 3.2(a), $2.75 \times 10^{-3} \leq \mu \leq 3.55 \times 10^{-3}$ and $0.145 \leq x \leq 0.163$.

$\frac{\partial f}{\partial \mu}(x_*; \mu_*) > 0$, then the map f_μ undergoes a backward saddle-node bifurcation (i.e., the node attractor is destroyed at x_* as μ increases through μ_*). If the inequality in either (c) or (d) is reversed, then the map undergoes a forward saddle-node bifurcation, while, if both these inequalities are reversed, the bifurcation remains backward. A saddle-node bifurcation in a one dimensional map is also called a tangent or a fold bifurcation.

3.2.1 Model

As an illustration of an indeterminate saddle-node bifurcation in a one-dimensional map, we construct an example in the following way. We consider the logistic map for a parameter value where there is a stable period three orbit. We denote this map $g(x)$

and its third iterate $g^{[3]}(x)$. The map $g^{[3]}(x)$ has three stable fixed points. We perturb the map $g^{[3]}(x)$ by adding a function (which depends on a parameter μ) that will cause a saddle-node bifurcation of one of the attracting fixed points but not of the other two [see Figs. 3.1(a) and 3.1(b)]. We investigate

$$f_\mu(x) = g^{[3]}(x) + \mu \sin(3\pi x), \quad \text{where} \quad g(x) = 3.832\, x(1-x). \tag{3.1}$$

Numerical calculations show that the function $f_\mu(x)$ satisfies all the conditions of the saddle-node bifurcation theorem for having a backward saddle-node bifurcation at $x_* \approx 0.15970$ and $\mu_* \approx 0.00279$. Figure 3.2(a) displays how the basins of the three attracting fixed points of the map f_μ change with variation of μ. For $\mu = 0$ the third iterate of the logistic map is unperturbed, and it has three attracting fixed points whose basins we color-coded with blue, green and red. For every value of μ, the red region $R[\mu]$ is the set of initial conditions attracted to the rightmost stable fixed point which we denote R_μ. The green region $G[\mu]$ is the set of initial conditions attracted to the middle stable fixed point which we denote G_μ. The blue region $B[\mu]$ is the set of initial conditions attracted to the leftmost stable fixed point which we denote B_μ.

For $\mu < \mu_*$, each of these colored sets has infinitely many disjoint intervals and a fractal boundary. As μ increases, the leftmost stable fixed point B_μ is destroyed via a saddle-node bifurcation on the fractal basin boundary. In fact, in this case, for $\mu < \mu_*$, every boundary point of one basin is a boundary point for all three basins. (That is, an arbitrarily small x-interval centered about any point on the boundary of any one of the basins contains pieces of the other two basins.) The basins are so-called Wada basins [31]. This phenomenon of a saddle-node bifurcation on the fractal boundary of Wada basins also occurs for the damped forced oscillators studied in Refs. [29, 30]. Alternatively, if we look at the saddle-node bifurcation as μ decreases through the value μ_*, then the basin $B[\mu]$ of the newly created stable fixed point immediately has

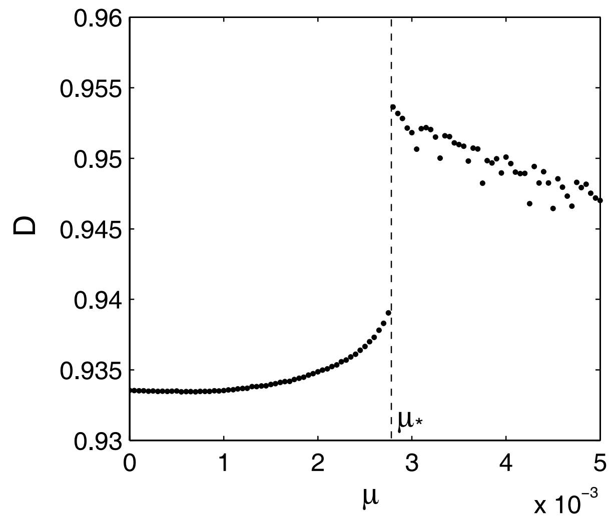

Figure 3.3: Fractal dimension of the basin boundary versus μ. Notice the continuous variation for $\mu < \mu_*$ and the discontinuous jump at μ_*, the parameter value at which the saddle-node bifurcation on the fractal basin boundary takes place.

infinitely many disjoint intervals and its boundary displays fractal structure. According to the terminology of Robert et al. [32], we may consider this bifurcation an example of an ‘explosion’.

3.2.2 Dimension of the Fractal Basin Boundary

Figure 3.3 graphs the computed dimension D of the fractal basin boundary versus the parameter μ. For $\mu < \mu_*$, we observe that D appears to be a continuous function of μ. Park et al. [33] argue that the fractal dimension of the basin boundary near μ_*, for $\mu < \mu_*$, scales as

$$D(\mu) \approx D_* - k(\mu_* - \mu)^{1/2}, \tag{3.2}$$

with D_* the dimension at $\mu = \mu_*$ (D_* is less than the dimension of the phase space), and k a positive constant. Figure 3.3 shows that the boundary dimension D experiences a discontinuous jump at the saddle-node bifurcation when $\mu = \mu_*$. We believe that this is due to the fact that the basin $B[\mu]$ suddenly disappears for $\mu > \mu_*$.

The existence of a fractal basin boundary has important practical consequences. In particular, for the purpose of determining which attractor eventually captures a given orbit, the arbitrarily fine-scaled structure of fractal basin boundaries implies considerable sensitivity to small errors in initial conditions. If we assume that initial points cannot be located more precisely than some $\epsilon > 0$, then we cannot determine which basin a point is in, if it is within ϵ of the basin boundary. Such points are called ϵ-uncertain. The Lebesgue measure of the set of ϵ-uncertain points (in a bounded region of interest) scales like $\epsilon^{D_0 - D}$, where D_0 is the dimension of the phase space ($D_0 = 1$ for one dimensional maps) and D is the box-counting dimension of the basin boundary [19]. For the case of a fractal basin boundary $(D_0 - D) < 1$. When $D_0 - D$ is small, a large decrease in ϵ results in a relatively small decrease in $\epsilon^{D_0 - D}$. This is discussed in Ref. [19] which defines the uncertainty dimension, D_u, as follows. Say we randomly pick an initial condition x with uniform probability density in a state-space region S. Then we randomly pick another initial condition y in S, such that $|y - x| < \epsilon$. Let $p(\epsilon, S)$ be the probability that x and y are in different basins. [We can think of $p(\epsilon, S)$ as the probability that an error will be made in determing the basin of an initial condition if the initial condition has uncertainty of size ϵ.] The uncertainty dimension of the basin boundary D_u is defined as the limit of $\ln p(\epsilon, S)/\ln(\epsilon)$ as ϵ goes to zero [19]. Thus, the probability of error scales as $p(\epsilon, S) \sim \epsilon^{D_0 - D_u}$, where for fractal basin boundaries $D_0 - D_u < 1$. This indicates enhanced sensitivity to small uncertainty in initial conditions. For example, if $D_0 - D_u = 0.2$, then a decrease of the initial condition uncertainty ϵ by a factor of 10 leads to only a relative small decrease in the final state uncertainty $p(\epsilon, S)$, since p decreases by a factor of about $10^{0.2} \approx 1.6$.

Thus, in practical terms, it may be essentially impossible to significantly reduce the final state uncertainty. In Ref. [19] it was conjectured that the box-counting dimension equals the uncertainty dimension for basin boundaries in typical dynamical systems. In Ref. [35] it is proven that the box-counting dimension, the uncertainty dimension and the Hausdorff dimension are all equal for the basin boundaries of one and two dimensional systems that are uniformly hyperbolic on their basin boundary.

We now explain some aspects of the character of the dependence of D on μ (see Fig. 3.3). From Refs. [36] it follows that the box-counting dimension and the Hausdorff dimension coincide for all intervals of μ for which the map f_μ is hyperbolic on the basin boundary, and that the dimension depends continuously on the parameter μ in these intervals. For $\mu > \mu_*$, there are many parameter values for which the map has a saddle-node bifurcation of a periodic orbit on the fractal basin boundary. At such parameter values, which we refer to as saddle-node bifurcation parameter values, the dimension is expected to be discontinuous (as it is at the saddle-node bifurcation of the fixed point, $\mu = \mu_*$, see Fig. 3.3). In fact, there exist sequences of saddle-node bifurcation parameter values converging to μ_* [34]. Furthermore, for each parameter value $\mu > \mu_*$ for which the map undergoes a saddle-node bifurcation, there exists a sequence of saddle-node bifurcation parameter values converging to that parameter value. The basins of attraction of the periodic orbits created by saddle-node bifurcations of high period exist only for very small intervals of the parameter μ. We did not encounter them numerically by iterating initial conditions for a discrete set of values of the parameter μ, as we did for the basin of our fixed point attractor.

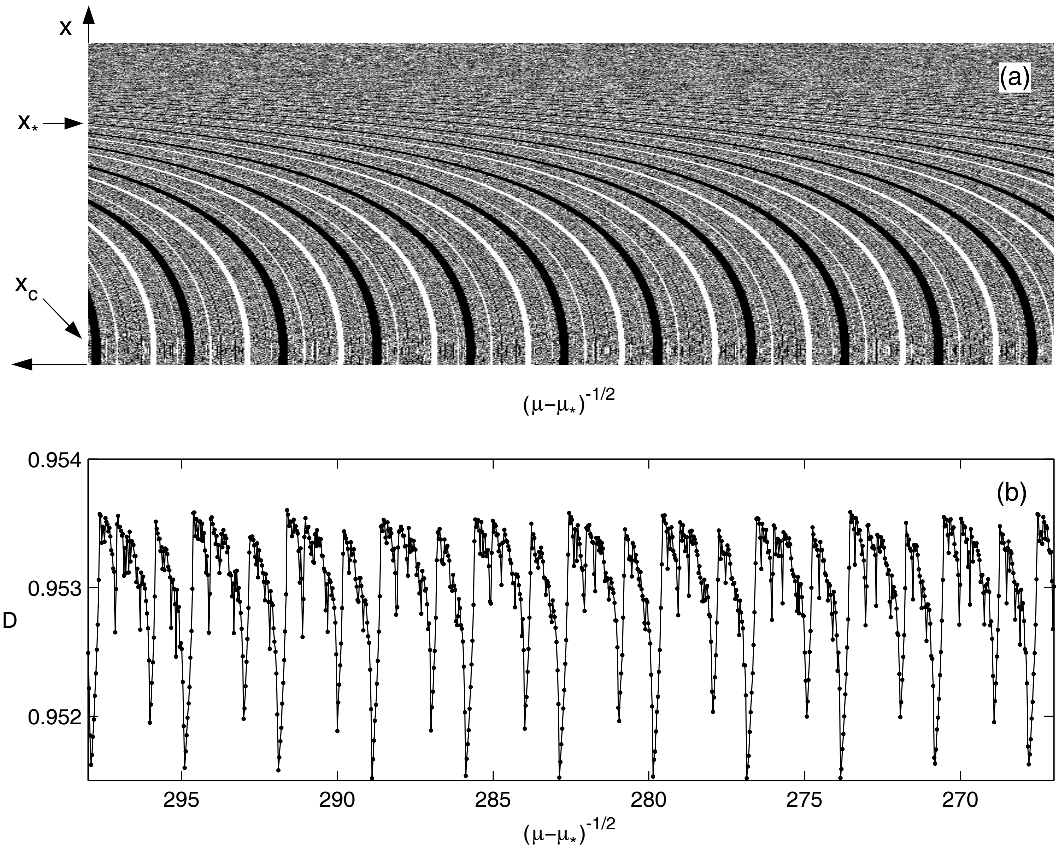

Figure 3.4: (a) Detail of Figure 3.2(b), with the horizontal axis changed from μ to $(\mu-\mu_*)^{-1/2}$ for $\mu > \mu_*$; $2.75\times10^{-3} \leq \mu \leq 3.55\times10^{-3}$ and $0.145 \leq x \leq 0.163$ The green stripes from Fig. 3.2(b) are colored black and the red stripes are colored white. The approximate position of the point x_* where the saddle-node bifurcation takes place is shown. x_c indicates the nearest critical point. (b) Detail of Fig. 3.3, displaying how the box dimension D of the fractal basin boundary varies with $1/(\mu_*-\mu)^{1/2}$. The horizontal axis of Figs. 3.4(a) and 3.4(b) are identical.

3.2.3 Scaling of the Fractal Basin Boundary

Just past μ_*, the remaining green and red basins display an alternating stripe structure [see Fig. 3.2(b)]. The red and green stripes are interlaced in a fractal structure. As we approach the bifurcation point, the interlacing becomes finer and finer scaled, with the scale approaching zero as μ approaches μ_*. Similar fine scaled structure is present in the neighborhood of all preiterates of x_*. If one changes the horizontal axis of Figs. 3.2(a,b) from μ to $(\mu - \mu_*)^{-1/2}$, then, the complex alternating stripe structure appears asymptotically periodic [see Fig. 3.4(a)]. [Thus, with identical horizontal scale, the dimension plot in Fig. 3.4(b) appears asymptotically periodic, as well.] We now explain why this is so. We restrict our discussion to a small neighborhood of x_*. Consider the second order expansion of f_μ in the vicinity of x_* and μ_*

$$\hat{f}_{\hat{\mu}}(\hat{x}) = \hat{\mu} + \hat{x} + a\hat{x}^2, \quad \text{where} \begin{cases} \hat{x} = x - x_*, \\ \hat{\mu} = \mu - \mu_*, \end{cases} \tag{3.3}$$

and $a \approx 89.4315$. The trajectories of $\hat{f}_{\hat{\mu}}$ in the neighborhood of $\hat{x} = 0$, for $\hat{\mu}$ close to zero, are good approximations to trajectories of f_μ in the neighborhood of $x = x_*$, for μ close to μ_*. Assume that we start with a certain initial condition for $\hat{f}_{\hat{\mu}}$, $\hat{x}_0 = \hat{x}_s$, and we ask the following question: What are all the positive values of the parameter $\hat{\mu}$ such that a trajectory passes through a fixed position $\hat{x}_f > 0$ at some iterate n? For any given x_f which is not on the fractal basin boundary, there exists a range of μ such that iterates of x_f under f_μ evolve to the same final attractor, for all values of μ in that range. In particular, once $a\hat{x}^2$ appreciably exceeds $\hat{\mu}$, the subsequent evolution is approximately independent of $\hat{\mu}$. Thus, we can choose $\hat{x}_f \gg \sqrt{\hat{\mu}/a}$, but still small enough so that it lies in the region of validity of the canonical form (3.3). There exists a range of such $\hat{x}_f$ values satisfying these requirements provided that $|\hat{\mu}|$ is small enough.

Since consecutive iterates of $\hat{f}_{\hat{\mu}}$ in the neighborhood of $\hat{x} = 0$ for $\hat{\mu}$ close to zero

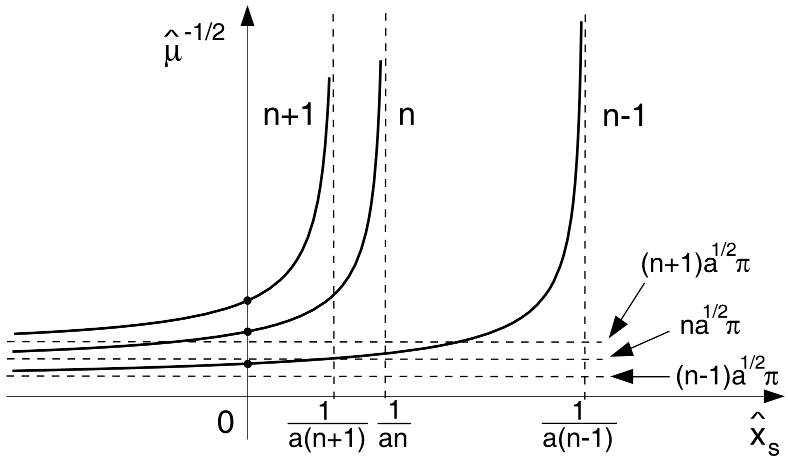

Figure 3.5: Qualitative graphs of the solution of Eq. (3.8), $\hat{\mu}_n^{-1/2}(\hat{x}_0)$, for three consecutive values of n. Note the horizontal asymptotes [$\hat{\mu}^{-1/2} = (n-1)a^{1/2}\pi$, $n\,a^{1/2}\pi$, and $(n+1)a^{1/2}\pi$], the vertical asymptotes [$\hat{x}_s = (a(n-1))^{-1}$, $(an)^{-1}$, and $(a(n+1))^{-1}$], both shown as dashed lines, and the intersections of the solid curves with $\hat{x}_0 = 0$ which are marked with black dots.

differ only slightly, we approximate the one dimensional map,

$$\hat{x}_{n+1} = \hat{f}_{\hat{\mu}}(\hat{x}_n) = \hat{\mu} + \hat{x}_n + a\hat{x}_n^2, \tag{3.4}$$

by the differential equation [27],

$$\frac{d\hat{x}}{dn} = \hat{\mu} + a\hat{x}^2, \tag{3.5}$$

where in (3.5) n is considered as a continuous, rather than a discrete, variable. Integrating (3.5) from $\hat{x}_s$ to $\hat{x}_f$ yields

$$n\sqrt{a\hat{\mu}} = \arctan\left(\sqrt{\frac{a}{\hat{\mu}}}\hat{x}_f\right) - \arctan\left(\sqrt{\frac{a}{\hat{\mu}}}\hat{x}_s\right). \tag{3.6}$$

Close to the saddle-node bifurcation (i.e., $0 < \hat{\mu} \ll 1$, and $\hat{x}_{s,f}$ close to zero), $\hat{f}_{\hat{\mu}}$ is a good approximation to f_μ. For $|\hat{x}_{s,f}|\sqrt{(a/\hat{\mu})} \gg 1$ Eq. (3.6) becomes

$$n\sqrt{a\hat{\mu}} \approx \pi. \tag{3.7}$$

The values of $\hat{\mu}_n^{-1/2}$ satisfying Eq. (3.7) increase with n in step of $\sqrt{a}/\pi$. For our example we have $a \approx 89.4315$, thus $\sqrt{a}/\pi \approx 3.010$. Counting many periods like those in Fig. 3.4 in the region of x_c, the closest critical point to x_* [see Fig. 3.4(a)], we find that the period of the stripe structure is 3.015, which is in good agreement with our theoretical value.

In order to investigate the structure of the fractal basin boundary in the vicinity of the saddle-node bifurcation (i.e., $\hat{x}_s$ close to $\hat{x}_* = 0$), we consider (3.6) in the case where we demand only $|\hat{x}_f|\sqrt{(a/\hat{\mu})} \gg 1$. Thus, Eq. (3.6) becomes

$$n\sqrt{a\hat{\mu}} \approx \frac{\pi}{2} - \arctan\left(\sqrt{\frac{a}{\hat{\mu}}}\hat{x}_s\right). \tag{3.8}$$

Let $\hat{\mu}_n^{-1/2}(\hat{x}_s)$ denote the solution of Eq. (3.8) for $\hat{\mu}$. Equation (3.8) implies the behavior of $\hat{\mu}_n^{-1/2}(\hat{x}_s)$ as function of $\hat{x}_s$ and n as sketched in Fig. 3.5. For a fixed n, $\hat{\mu}_n^{-1/2}$ has a horizontal asymptote at the value $n\sqrt{a}/\pi$ as $\hat{x}_s \to -\infty$, and a vertical asymptote to infinity at $\hat{x}_s = 1/(an)$. For $\hat{x}_s < 0$, we have an infinite number of values of the parameter $\hat{\mu}$, for which an orbit of $\hat{f}_{\hat{\mu}}$ starting at $\hat{x}_s$ passes through the same position $\hat{x}_f$, after some number of iterations. For $\hat{x}_s = 0$ (i.e., $x_s = x_*$), we also have an infinite number of $\hat{\mu}_n^{-1/2}(0)$, but with constant step $2\sqrt{a}/\pi$ rather than $\sqrt{a}/\pi$ (see the intersections marked with black dots in the Fig. 3.5). This is hard to verify from numerics, since $\frac{\partial\hat{\mu}_n^{-1/2}}{\partial\hat{x}_s}(0) = a^{3/2}(2n/\pi)^2$ increases with n^2, and the stripes become very tilted in the neighborhood of $\hat{x}_s = \hat{x}_* = 0$. [See Fig. 3.4(a), where the approximate positions of x_c and x_* on the vertical axis are indicated.] For $\hat{x}_s > 0$, $\hat{\mu}_n^{-1/2}$ has only a limited number of values with $n_{\max} < 1/(a\hat{x}_0)$.

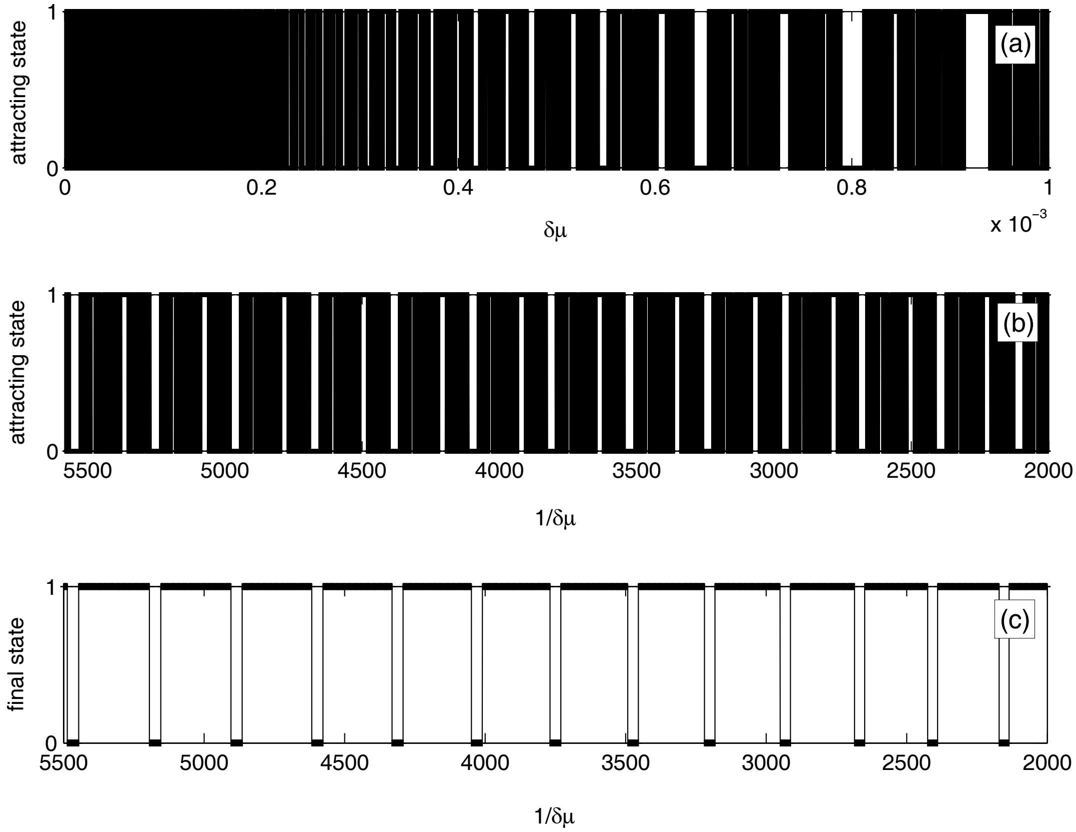

Figure 3.6: (a) Final attracting state of swept orbits versus $\delta\mu$. We have chosen $\mu_s = \hat{\mu}_s + \mu_* = 0$, and $\mu_f = 4.5 \times 10^{-3}$. The attractor R_{μ_f} is represented by 1 and the attractor G_{μ_f} is represented by 0. (b) Detail of Fig. 3.6(b) with the horizontal scale changed from $\delta\mu$ to $1/\delta\mu$. The structure of white and black bands becomes asymptotically periodic. (c) Final state of orbits for the system $\hat{f}_{\hat{\mu}}$ versus $1/\delta\mu$. The final state of an orbit is defined to be 0 if there exists n such that $100 < \hat{x}_n < 250$, and is defined to be 1, otherwise. We have chosen $\hat{\mu}_s = -\mu_*$, so that Figs. 3.6(b,c) have the same asymptotic periodicity.

3.2.4 Sweeping Through an Indeterminate Saddle-Node Bifurcation

In order to understand the consequences of a saddle-node bifurcation on a fractal basin boundary for systems experiencing slow drift, we imagine the following experiment. We start with the dynamical system f_μ at parameter $\mu_s < \mu_*$, with x_0 on the attractor to be destroyed at $\mu = \mu_*$ by a saddle-node bifurcation (i.e., B_μ). Then, as we iterate, we slowly change μ by a small constant amount $\delta\mu$ per iterate, thus increasing μ from μ_s to $\mu_f > \mu_*$,

$$\begin{aligned} x_{n+1} &= f_{\mu_n}(x_n), \\ \mu_n &= \mu_s + n\,\delta\mu. \end{aligned} \tag{3.9}$$

When $\mu \geq \mu_f$ we stop sweeping the parameter μ, and, by iterating further, we determine to which of the remaining attractors of f_{μ_f} the orbit goes. Numerically, we observe that, if $(\mu_f - \mu_*)$ is not too small, then, by the time μ_f is reached, the orbit is close to the attractor of f_{μ_f} to which it goes. [From our subsequent analysis, 'not too small $|\mu_{s,f} - \mu_*|$' translates to choices of $\delta\mu$ that satisfy $(\delta\mu)^{2/3} \ll |\mu_{s,f} - \mu_*|$.] We repeat this for different values of $\delta\mu$ and we graph the final attractor position for the orbit versus $\delta\mu$ [see Fig. 3.6(a)]. For convenience in the graphical representation of Figs. 3.6(a,b), we have represented the attractor of the green region $G[\mu]$, denoted G_{μ_f}, as a 0, and the attractor of the red region $R[\mu]$, denoted R_{μ_f}, as a 1. In Fig. 3.6(a) we use of 25,000 points having the vertical coordinate either 0 or 1, which we connect with straight lines. In an interval of $\delta\mu$ for which the system reaches the same final attractor (either 0 or 1), the lines connecting the points are horizontal. Such intervals appear as white bands in Fig. 3.6, if they are wider than the width of the plotted lines connecting 0's and 1's. For example, in Fig. 3.6(a), the white band centered at

$\delta\mu = 0.8 \times 10^{-3}$ has at the bottom a thick horizontal line, which indicates that for the whole of that interval, the orbit reaches the attractor G_{μ_f} which we represented by 0. Adjacent intervals of width less than the plotted lines appear as black bands. Within such black bands, an uncertainty in $\delta\mu$ of size equal to the width of the plotted line makes the attractor that the orbit goes to indeterminate. Figure 3.6(a) shows that the widths of the white bands decrease as $\delta\mu$ decreases, such that, for small $\delta\mu$, we see only black.

If $(\mu_f - \mu_*)$ is large enough (i.e., $(\delta\mu)^{2/3} \ll |\mu_f - \mu_*|$), numerics and our subsequent analysis show that Fig. 3.6 is independent of μ_f. This fact can be understood as follows. Once $\mu = \mu_f$, the orbit typically lands in the green or the red basin of attraction and goes to the corresponding attractor. Due to sweeping, it is possible for the orbit to switch from being in one basin of attraction of the *time-independent* map f_μ to the other, since the basin boundary between $G[\mu]$ and $R[\mu]$ changes with μ. However, the sweeping of μ is slow (i.e., $\delta\mu$ is small), and, once $(\mu - \mu_*)$ is large enough, the orbit is far enough from the fractal basin boundary, and the fractal basin boundary changes too little to switch the orbit between $G[\mu]$ and $R[\mu]$.

We also find numerically that Figs. 3.6(a,b) are independent of the initial condition x_0, provided that it is in the blue basin $B[\mu_s]$, sufficiently far from the fractal basin boundary, and that $|\mu_s - \mu_*|$ is not too small (i.e., $(\delta\mu)^{2/3} \ll |\mu_s - \mu_*|$).

If one changes the horizontal scale of Fig. 3.6(a) from $\delta\mu$ to $1/\delta\mu$ [see Fig. 3.6(b)], the complex band structure appears asymptotically periodic. Furthermore, we find that the period in $(1/\delta\mu)$ of the structure in Fig. 3.6(b) asymptotically approaches $-1/(\mu_s - \mu_*)$ as $\delta\mu$ becomes small.

In order to explain this result, we again consider the map $\hat{f}_{\hat{\mu}}$, the local approximation of f_μ in the region of the saddle-node bifurcation. Equations (3.9) can be approximated by

$$\begin{aligned} \hat{x}_{n+1} &= \hat{f}_{\hat{\mu}_n}(\hat{x}_n) = \hat{\mu}_n + \hat{x}_n + a\hat{x}_n^2, \\ \hat{\mu}_n &= \hat{\mu}_s + n\,\delta\mu. \end{aligned} \tag{3.10}$$

We perform the following numerical experiment. We consider orbits of our approximate two dimensional map given by Eq. (3.10) starting at $\hat{x}_s = -\sqrt{-\hat{\mu}_s/a}$. We define a final state function of an orbit swept with parameter $\delta\mu$ in the following way. It is 0 if the orbit has at least one iterate in a specified fixed interval far from the saddle-node bifurcation, and is 1, otherwise. In particular, we take the final state of a swept orbit to be 0 if there exists n such that $100 < \hat{x}_n < 250$, and to be 1 otherwise. Figure 3.6(c) graphs the corresponding numerical results. Similar to Fig. 3.6(b), we observe periodic behavior in $1/\delta\mu$ with period $-1/\hat{\mu}_s$. In contrast to Fig. 3.6(b) where the white band structure seems fractal, the structure within each period in Fig. 3.6(c) consists of only one interval where the final state is 0 and one interval where the final state is 1. This is because $100 < \hat{x} < 250$ is a single interval, while the green basin [denoted 0 in Fig. 3.6(b)] has an infinite number of disjoint intervals and a fractal boundary (see Fig. 3.2).

With the similarity between Figs. 3.6(b) and 3.6(c) as a guide, we are now in a position to give a theoretical analysis explaining the observed periodicity in $1/\delta\mu$. In particular, we now know that this can be explained using the canonical map (3.10), and that the periodicity result is thus universal [i.e., independent of the details of our particular example, Eq. (3.1)]. For slow sweeping (i.e., $\delta\mu$ small), consecutive iterates of (3.10) in the vicinity of $\hat{x} = 0$ and $\hat{\mu} = 0$ differ only slightly, and we further approximate the system by the following Ricatti differential equation,

$$\frac{d\hat{x}}{dn} = \hat{\mu}_s + n\delta\mu + a\hat{x}^2. \tag{3.11}$$

The solution of Eq. (3.11) can be expressed in terms of the Airy functions Ai and Bi and their derivatives, denoted by Ai' and Bi',

$$\hat{x}(n) = \frac{\eta Ai'(\xi) + Bi'(\xi)}{\eta Ai(\xi) + Bi(\xi)} \left(\frac{\delta\mu}{a^2}\right)^{1/3}, \tag{3.12}$$

where

$$\xi(n) = -a^{1/3}\,\frac{\hat{\mu}_s + n\,\delta\mu}{\delta\mu^{2/3}}, \tag{3.13}$$

and η is a constant to be determined from the initial condition. We are only interested in the case of slow sweeping, $\delta\mu \ll 1$, and $\hat{x}(0) \equiv \hat{x}_s = -\sqrt{-\hat{\mu}_s/a}$ (which is the stable fixed point of $\hat{f}_{\hat{\mu}}$ destroyed by the saddle-node bifurcation at $\hat{\mu} = 0$). In particular, we will consider the case where $\hat{\mu}_s < 0$ and $|\hat{\mu}_s| \gg \delta\mu^{2/3}$ (i.e., $|\xi(0)| \gg 1$). Using $\hat{x}(0) = -\sqrt{-\hat{\mu}_s/a}$ to solve for η yields $\eta \sim \mathcal{O}[\xi(0)e^{2\xi(0)}] \gg 1$. For positive large values of $\xi(n)$ (i.e., for n small enough), using the corresponding asymptotic expansions of the Airy functions [37], the lowest order in $\delta\mu$ approximation to (3.12) is

$$\hat{x}(n) \approx -\sqrt{-\frac{\hat{\mu}_s + n\,\delta\mu}{a}}, \tag{3.14}$$

with the correction term of higher order in $\delta\mu$ being negative. Thus, for n sufficiently smaller than $-\hat{\mu}_s/\delta\mu$, the swept orbit lags closely behind the fixed point for $\hat{f}_{\hat{\mu}}$ with $\hat{\mu}$ constant. For $\xi \leq 0$, we use the fact that η is large to approximate (3.12) as

$$\hat{x}(n) \approx \frac{Ai'(\xi)}{Ai(\xi)} \left(\frac{\delta\mu}{a^2}\right)^{1/3}. \tag{3.15}$$

Note that

$$\hat{x}(-\hat{\mu}_s/\delta\mu) \approx \frac{Ai'(0)}{Ai(0)} \left(\frac{\delta\mu}{a^2}\right)^{1/3} = (-0.7290...) \left(\frac{\delta\mu}{a^2}\right)^{1/3} \tag{3.16}$$

gives the lag of the swept orbit relative to the fixed point attractor evaluated at the saddle-node bifurcation. Equation (3.15) does not apply for $n > n_{\max}$, where $n_{\max}$ is the value of n for which $\xi(n_{\max}) = \tilde{\xi}$, the largest root of $Ai(\tilde{\xi}) = 0$ (i.e., $\tilde{\xi} = -2.3381...$). At $n = n_{\max}$, the normal form approximation predicts that the orbit diverges to $+\infty$. Thus, for n near $n_{\max}$, the normal form approximation of the dynamical system ceases to be valid. Note, however, that (3.15) can be valid even for $\xi(n)$ close to $\xi(n_{\max})$. This is possible because $\delta\mu$ is small. In particular, we can consider times up to the time n' where n' is determined by $\xi' \equiv \xi(n') = \tilde{\xi} + \delta\xi$, ($\delta\xi > 0$ is small,) provided $|\hat{x}(n')| \ll 1$ so that the normal form applies. That is, we require $[Ai'(\xi')/Ai(\xi')](\delta\mu/a^2)^{1/3} \ll 1$, which can be satisfied even if $[Ai'(\xi')/Ai(\xi')]$ is large. Furthermore, we will take the small quantity $\delta\xi$ to be not too small (i.e., $\delta\xi/(a\,\delta\mu)^{1/3} \gg 1$), so that $(n_{\max} - n') \gg 1$. We then consider (3.15) in the range, $-(\hat{\mu}_s/\delta\mu) \leq n < n'$, where the normal form is still valid.

We use Eq. (3.15) for answering the following question: What are all the values of the parameter $\delta\mu$ ($\delta\mu$ small) for which an orbit passes exactly through the same position $\hat{x}_f > 0$, at some iterate n_f? All such orbits would further evolve to the same final attractor, independent of $\delta\mu$, provided $a\hat{x}_f^2 \gg \hat{\mu}_s + n_f\,\delta\mu$; i.e., $\hat{x}_f$ is large enough that $\hat{\mu}_f = \hat{\mu}_s + n_f\,\delta\mu$ does not much influence the orbit after $\hat{x}$ reaches $\hat{x}_f$. [Denote $\xi(n_f)$ as $\xi(n_f) \equiv \xi_f$.] Using (3.15) we can estimate when this occurs, $a\hat{x}_f^2 = [Ai'(\xi_f)/Ai(\xi_f)]^2(\delta\mu^2/a)^{1/3} \gg (\hat{\mu}_s + n_f\,\delta\mu)$ or $[Ai'(\xi_f)/Ai(\xi_f)]^2 \gg \xi_f$. This inequality is satisfied when ξ_f gets near $\tilde{\xi}$, which is the largest zero of Ai (i.e., $\xi_f = \tilde{\xi} + \delta\xi$, where $\delta\xi$ is a small positive quantity). We now rewrite Eq. (3.15) in the following way

$$\frac{1}{\delta\mu} = -\frac{n_f}{\hat{\mu}_s - \left[\frac{(\delta\mu)^2}{a}\right]^{1/3} K\left[\left(\frac{a^2}{\delta\mu}\right)^{1/3}\hat{x}_f\right]}, \tag{3.17}$$

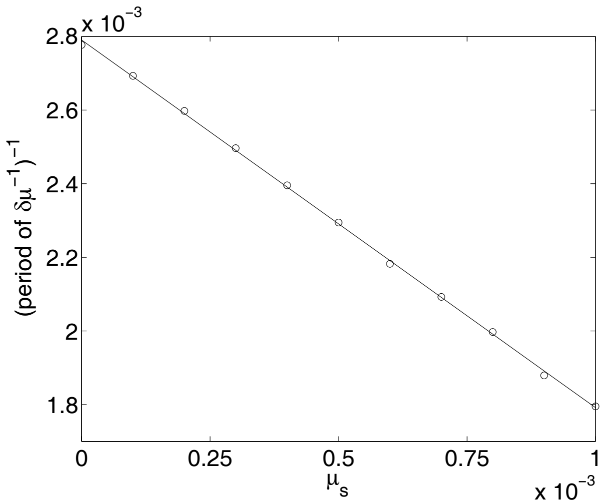

Figure 3.7: Numerical results for the inverse of the limit period in $1/\delta\mu$ versus μ_s. The fit line is $[\Delta\,(1/\delta\mu)]^{-1} = -0.9986\mu_s + 0.0028$ and indicates good agreement with the theoretical explanation presented in text.

representing a transcedental equation in $\delta\mu$ where $\hat{\mu}_s$ and $\hat{x}_f$ are fixed, n_f is a large positive integer (i.e., $n_f - 1$ is the integer part of $(\hat{\mu}_f - \hat{\mu}_s)/\delta\mu$), and $K(\zeta)$ is the inverse function of $Ai'(\xi)/Ai(\xi)$ in the neighborhood of $\zeta = (a^2/\delta\mu)^{1/3}\,\hat{x}_f \gg 1$. Thus $|K[(a^2/\delta\mu)^{1/3}\,\hat{x}_f]| \lesssim |K(\infty)| = |\tilde{\xi}|$. The difference $[1/\delta\mu(x_f, n_f{+}1) - 1/\delta\mu(x_f, n_f)]$, where $\delta\mu(x_f, n_f)$ is the solution of Eq. (3.17), yields the limit period of the attracting state versus $1/\delta\mu$ graph (see Fig. 3.6). We denote this limit period by $\Delta\,(1/\delta\mu)$. For small $\delta\mu$, the term involving $K[(a^2/\delta\mu)^{1/3}\,\hat{x}_f]$ in Eq. (3.17) can be neglected, and we get $\Delta\,(1/\delta\mu) = -\hat{\mu}_s^{-1} = (-\mu_s + \mu_*)^{-1}$. Figure 3.7 graphs numerical results of $[\Delta\,(1/\delta\mu)]^{-1}$ versus μ_s for our map example given by Eq. (3.9). The fit line is $[\Delta\,(1/\delta\mu)]^{-1} = -0.9986\mu_s + 0.0028$, which agrees well with the prediction of the above analysis and our numerical value for μ at the bifurcation, $\mu_* \approx 0.00279$.

An alternate point of view on this scaling property is as follows. For $\hat{\mu} < 0$ (i.e., $\mu < \mu_*$) and slow sweeping (i.e., $\delta\mu$ small), the orbit closely follows the sta-

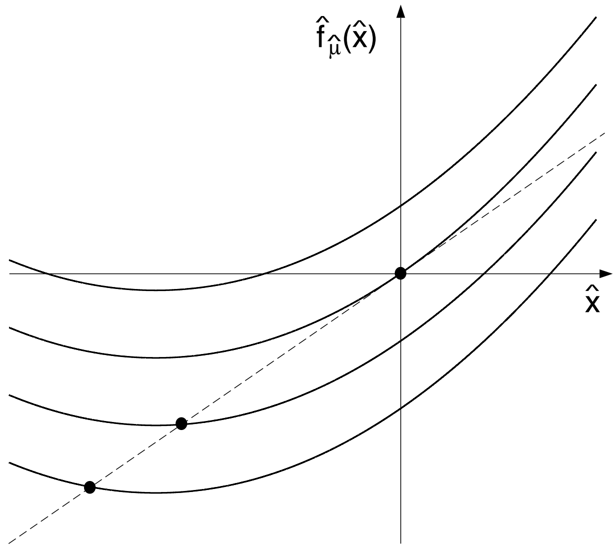

Figure 3.8: Graphs of $\hat{f}_{\hat{\mu}}(\hat{x})$ at different values of the parameter $\hat{\mu}$. The black dots indicate the stable fixed points of $\hat{f}_{\hat{\mu}}$ for different values of $\hat{\mu}$.

ble fixed point attractor of $\hat{f}_{\hat{\mu}}$, until $\hat{\mu} \geq 0$, and the saddle-node bifurcation takes place. However, due to the discreteness of n, the first nonnegative value of $\hat{\mu}$ depends on $\hat{\mu}_s$ and $\delta\mu$ (see Fig. 3.8). Now consider two values of $\delta\mu$, one $\delta\mu_m$ satisfying $\hat{\mu}_s + m\,\delta\mu_m = 0$, and another $\delta\mu_{m+1}$ satisfying $\hat{\mu}_s + (m+1)\,\delta\mu_{m+1} = 0$. Because $\delta\mu_m$ and $\delta\mu_{m+1}$ are very close (for large m) *and* both lead $\hat{\mu}(n)$ to pass through $\hat{\mu} = \hat{\mu}_* = 0$ (one at time $n = m$, and the other at time $n = m+1$), it is reasonable to assume that their orbits for $\hat{\mu}_s/\delta\mu < n < n'$ are similar (except for a time shift $n \to n+1$); i.e., they go to the same attractor. Thus, the period of $1/\delta\mu$ is approximately $\Delta\,(1/\delta\mu) = 1/\delta\mu_{m+1} - 1/\delta\mu_m = -\hat{\mu}_s^{-1}$.

We now consider the intervals of $1/\delta\mu$ between the centers of consecutive wide white bands in Fig. 3.6(b). Figure 3.9 graphs the calculated fractal dimension D' of the boundary between white bands in these consecutive intervals versus their center

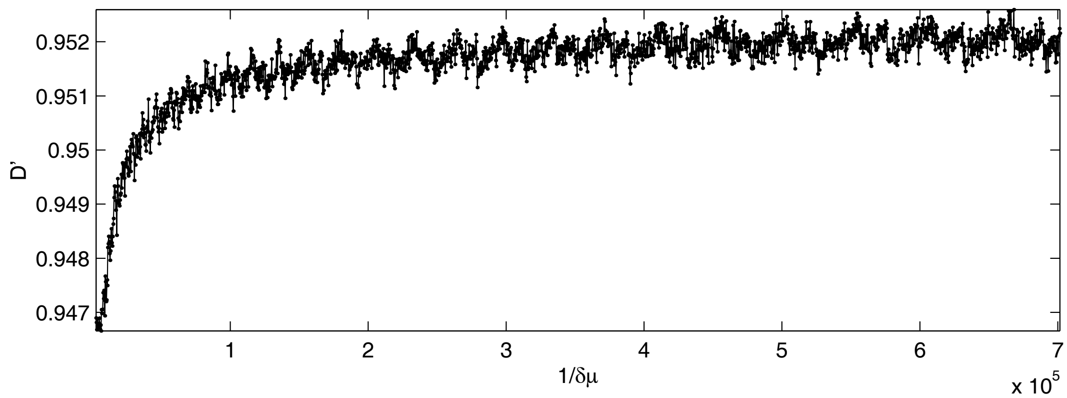

Figure 3.9: The calculated fractal dimension D' of the structure in the intervals between the centers of consecutive wide white bands in Fig. 3.6(b) versus their center value of $1/\delta\mu$.

value of $1/\delta\mu$. From Fig. 3.9, we see that as $1/\delta\mu$ increases, the graph of the fractal dimension D' does not converge to a definite value, but displays further structure. Nevertheless, numerics show that as $1/\delta\mu$ becomes large (i.e., in the range of 6.5×10^5), D' varies around the value 0.952. This is consistent with the numerics presented in Fig. 3.4(b) which graphs the dimension of the fractal basin boundary for the time-independent map f_μ, at fixed values of the parameter μ where $\mu > \mu_*$. Thus, for large $1/\delta\mu$, D' provides an estimate of the dimension of the fractal basin boundary in the absence of sweeping at $\mu > \mu_*$.

We now discuss a possible experimental application of our analysis. The conceptually most straightforward method of measuring a fractal basin boundary would be to repeat many experiments each with precisely chosen initial conditions. By determining the final attractor corresponding to each initial condition, basins of attraction could conceivably be mapped out [21]. However, it is commonly the case that accurate control of initial conditions is not feasible for experiments. Thus, the application of this direct method is limited, and, as a consequence, fractal basin boundaries have received little experimental study, in spite of their fundamental importance. If a saddle-node bifurcation occurs on the fractal basin boundary, an experiment can be arranged to take advantage of this. In this case, the purpose of the experiment would be to measure the dimension D' as an estimate of the fractal dimension of the basin boundary D. The measurements would determine the final attractor of orbits starting at the attractor to be destroyed by the saddle-node bifurcation, and swept through the saddle-node bifurcation at different velocities (i.e., the experimental data corresponding to the numerics in Fig. 3.6). This does not require precise control of the initial conditions of the orbits. It is sufficient for the initial condition to be in the basin of the attractor to be destroyed by the saddle-node bifurcation; after enough time, the orbit will be as close to the attractor as the noise level allows. Then, the orbit may be swept through the saddle-node bifurcation. The final states of the orbits are attractors; in their final states, orbits are robust to noise and to measurement perturbations. The only parameters which require rigorous control are the sweeping velocity (i.e., $\delta\mu$) and the initial value of the parameter to be swept (i.e., μ_s); precise knowledge of the parameter value where the saddle-node bifurcation takes place (i.e., μ_*) is not needed. [It is also required that the noise level be sufficiently low (see Sec. 3.2.6).]

3.2.5 Capture Time

A question of interest is how much time it takes for a swept orbit to reach the final attracting state. Namely, we ask how many iterations with $\mu > \mu_*$ are needed for the orbit to reach a neighborhood of the attractor having the green basin. Due to slow sweeping, the location of the attractor changes slightly on every iterate. If x_μ is a fixed

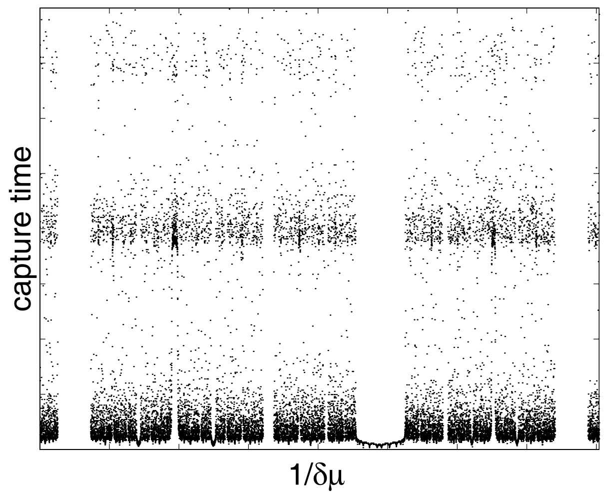

Figure 3.10: Capture time by the fixed point attractor G_{μ_f} versus $1/\delta\mu$. We have chosen $\mu_s = 0$. The range of $1/\delta\mu$ is approximatelly one period of the graph in Fig. 3.6(b), with $\delta\mu \approx 10^{-8}$. The vertical axis ranges between 250 and 650. No points are plotted for values of $\delta\mu$ for which the orbit reaches the fixed point attractor R_{μ_f}.

point attractor of f_μ (with μ constant), then a small change $\delta\mu$ in the parameter μ, yields a change in the position of the fixed point attractor,

$$(x_{\mu+\delta\mu} - x_\mu) \equiv \delta x = \delta\mu \, \frac{\frac{\partial f}{\partial \mu}(x_\mu;\mu)}{1 - \frac{\partial f_\mu}{\partial x}(x_\mu)}.$$

We consider the swept orbit to have reached its final attractor if consecutive iterates differ by about δx (which is proportional to $\delta\mu$). For numerical purposes, we consider that the orbit has reached its final state if $|x_{n+1} - x_n| < 10\,\delta\mu$. In our numerical experiments, this condition is satisfied by every orbit before μ reaches its final value μ_f. We refer to the number of iterations with $\mu > \mu_*$ needed to reach the final state as the *capture time* of the corresponding orbit. Figure 3.10 plots the capture time by the attractor G_{μ_f} [having the green basin in Fig. 3.2] versus $1/\delta\mu$ for a range corresponding to one period of the structure in Fig. 3.6(b). No points are plotted for values of

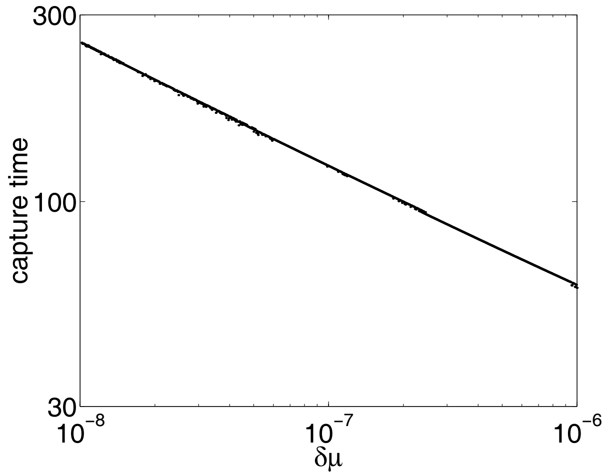

Figure 3.11: Capture time by the middle fixed point attractor of f_μ versus $\delta\mu$ ($\mu_s = 0$). The best fitting line (not shown) has slope -0.31, in agreement with the theory.

$\delta\mu$ for which the orbit reaches the attractor R_{μ_f}. The capture time graph has fractal features, since for many values of $\delta\mu$ the orbit gets close to the fractal boundary between $R[\mu]$ and $G[\mu]$. Using the fact that the final destination of the orbit versus $1/\delta\mu$ is asymptotically periodic [see Fig. 3.6(b)], we can provide a further description of the capture time graph. We consider the series of the largest intervals of $1/\delta\mu$ for which the orbit reaches the attractor G_{μ_f} [see Fig. 3.6(b); we refer to the wide white band around $1/\delta\mu = 2400$ and the similar ones which are (asymptotically) separated by an integer number of periods]. Orbits swept with $\delta\mu$ at the centers of these intervals spend only a small number of iterations close to the common fractal boundary of $R[\mu]$ and $G[\mu]$. Thus, the capture time of such similar orbits does not depend on the structure of the fractal basin boundary. We use Eq. (3.15) as an approximate description of these orbits. A swept orbit reaches its final attracting state as $\hat{x}(n)$ becomes large. Then, the orbit is rapidly trapped in the neighborhood of one of the swept attractors of f_μ. Thus, we equate the argument of the Airy function in the denominator to its first root [see (3.15)], solve for n, and substract $-\hat{\mu}_s/\delta\mu$ (the time for $\hat{\mu}$ to reach the bifurcation value). This yields the following approximate formula for the capture time

$$n_C \approx |\tilde{\xi}|(a\,\delta\mu)^{-1/3}, \tag{3.18}$$

where $\tilde{\xi} = -2.3381...$ is the largest root of the Airy function Ai. Thus, we predict that for small $\delta\mu$, a log-log plot of the capture time of the selected orbits versus $\delta\mu$ is a straight line with slope -1/3. Figure 3.11 shows the corresponding numerical results. The best fitting line (not shown) has slope -0.31, in agreement with our prediction [38].

3.2.6 Sweeping Through an Indeterminate Saddle-Node Bifurcation in the Presence of Noise

We now consider the addition of noise. Thus, we change our swept dynamical system to

$$\begin{aligned} x_{n+1} &= f_{\mu_n}(x_n) + A\,\epsilon_n, \\ \mu_n &= \mu_s + n\,\delta\mu, \end{aligned} \tag{3.19}$$

where ϵ_n is random with uniform probability density in the interval $[-1, 1]$, and A is a parameter which we call the noise amplitude. See Fig. 3.6(a) which shows the numerical results of the final destination of the orbits versus $\delta\mu$ in the case $A = 0$. The graph exhibits fractal features of structure at arbitrarily small scales. The addition of small noise is expected to alter this structure, switching the final destination of orbits. In this case, it is appropriate to study the probability of orbits reaching one of the final destinations. For every A, we compute the final attractor of a large number of orbits having identical initial condition and parameters, but with different realizations of the noise. We estimate the probability that an orbit reaches a certain attractor by

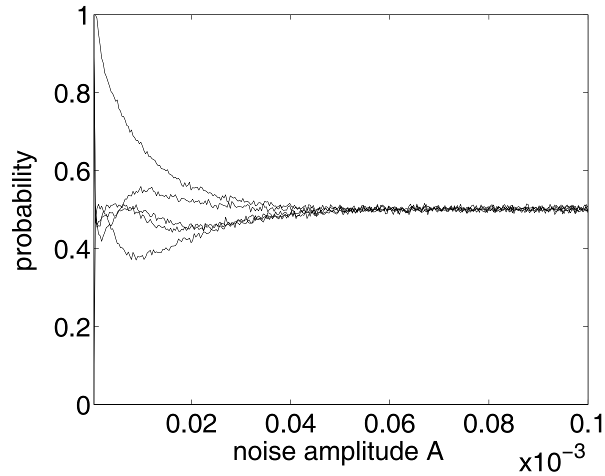

Figure 3.12: Probability that one orbit reaches the middle fixed point attractor of f_μ versus the noise amplitude A, for five different values of $\delta\mu$ (10^{-5}, $10^{-5} \pm 2.5 \times 10^{-8}$ and $10^{-5} \pm 5 \times 10^{-8}$). We have chosen $\mu_s = 0$.

the fraction of such orbits that have reached the specified attractor in our numerical simulation. Figure 3.12 graphs the probability that an orbit reaches the attractor G_{μ_f} versus the noise amplitude A. We present five graphs corresponding to five different values of $\delta\mu$ equally spaced in a range of 10^{-7} centered at 10^{-5} (i.e., $\delta\mu = 10^{-5}$, $10^{-5} \pm 2.5 \times 10^{-8}$ and $10^{-5} \pm 5 \times 10^{-8}$). We notice that the probability graphs have different shapes, but a common horizontal asymptote in the limit of large noise. The value of the horizontal asymptote, approximately equal to 0.5, is related to the relative measure of the corresponding basin.

As in the previous subsection, we take advantage of the asymptotically periodic structure of the noiseless final destination graph versus $1/\delta\mu$ [see Fig. 3.6(b)]. We consider centers of the largest intervals of $1/\delta\mu$ for which an orbit reaches the middle attractor in the absence of noise. We chose five such values of $\delta\mu$, spread over two decades, where the ratio of consecutive values is approximately 3. Figure 3.13(a)

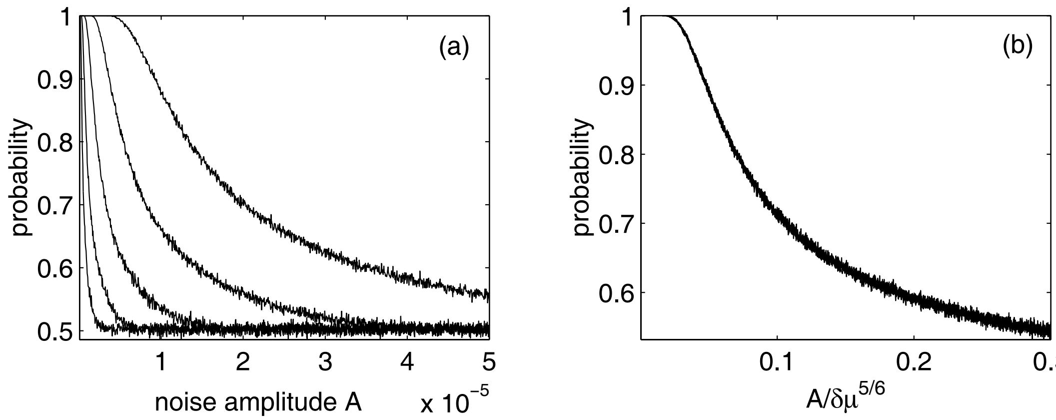

Figure 3.13: Probability that an orbit reaches the middle fixed point attractor of f_μ, for five selected values of $\delta\mu$ spread over two decades: (a) versus the noise amplitude A, and (b) versus $A/(\delta\mu)^{5/6}$, We have chosen $\mu_s = 0$.

graphs the probability that an orbit reaches the middle fixed point attractor versus the noise amplitude A, for the five selected values of $\delta\mu$. From right to left, the $\delta\mu$ values corresponding to the curves are approximately: 3.445974×10^{-5}, 1.147767×10^{-5}, 3.820744×10^{-6}, 1.273160×10^{-6} and 4.243522×10^{-7}. We notice that all the curves have qualitatively similar shape. For a range from zero to small A, the probability is 1, and as A increases, the probability decreases to a horizontal asymptote. The rightmost curve in the family corresponds to the largest value of $\delta\mu$ ($\delta\mu \approx 3.445974 \times 10^{-5}$), and the leftmost curve corresponds to the smallest value of $\delta\mu$ ($\delta\mu \approx 4.243522 \times 10^{-7}$). Figure 3.13(b) shows the same family of curves as in Fig. 3.13(a), but with the horizontal scale changed from A to $A/(\delta\mu)^{5/6}$. All data collapse to a single curve, indicating that the probability that a swept orbit reaches the attractor G_{μ_f} depends only on the reduced variable $A/(\delta\mu)^{5/6}$. Later, we provide a theoretical argument for this scaling.

In order to gain some understanding of this result, we follow the idea of Sec. 3.2.4,

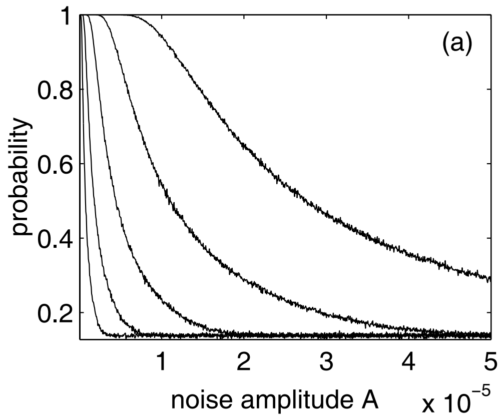

Figure 3.14: Probability that an orbit of $\hat{f}_{\hat{\mu}}$ reaches a fixed interval far from the saddle-node bifurcation (i.e., [100, 250]), for five values of $\delta\mu$ spread over two decades: (a) versus the noise amplitude A, and (b) versus $A/(\delta\mu)^{5/6}$. We have chosen $\mu_s = 0$.

and use the canonical form $\hat{f}_{\hat{\mu}}$ to propose a simplified setup of our problem. We modify (3.10) by the addition of a noise term $A\,\epsilon_n$ in the right hand side of the first equation of (3.10). We are interested in the probability that a swept orbit has at least one iterate, $\hat{x}_n$, in a specified fixed interval far from the vicinity of the saddle-node bifurcation. More precisely, we analyze how this probability changes versus A and $\delta\mu$. Depending on the choice of interval and the choice of $\delta\mu$, the probability versus A graph (not shown) has various shapes. For numerical purposes, we choose our fixed interval to be the same as that of Sec. 3.2.4, $100 \leq \hat{x} \leq 250$. We then select values of $\delta\mu$ for which a noiseless swept orbit, starting at $\hat{x}_s = -\sqrt{-\hat{\mu}_s/a}$, reaches exactly the center of our fixed interval. The inverse of these values of $\delta\mu$ are centers of intervals where the final state of the swept orbits is 0 [see Fig. 3.6(c)]. We consider five such values of $\delta\mu$, where the ratio of consecutive values is approximately 3. Figure 3.14(a) shows the probability that a swept orbit has an iterate in our fixed interval versus the noise amplitude for the selected values of $\delta\mu$. From right to left, the $\delta\mu$ values corresponding to the curves are approximately: 3.451540$\times 10^{-5}$, 1.149162$\times 10^{-5}$, 3.829769$\times 10^{-6}$, 1.276061$\times 10^{-6}$ and 4.253018$\times 10^{-7}$. Figure 3.14(a) shares the qualitative characteristics of Fig. 3.13(a), with the only noticeable difference that the value of the horizontal asymptote is now approximately 0.1. Figure 3.14(b) shows the same family of curves as in Fig. 3.14(a), where the horizontal scale has been changed from A to $A/(\delta\mu)^{5/6}$. As for Fig. 3.12(b), this achieves good collapse of the family of curves.

We now present a theoretical argument for why the probability of reaching an attractor depends on $\delta\mu$ and A only through the scaled variable $A/(\delta\mu)^{5/6}$ when $\delta\mu$ and A are small. From our results in Figs. 3.14, we know that the scaling we wish to demonstrate should be obtainable by use of the canonical form $\hat{f}_{\hat{\mu}}$. Accordingly, we again use the differential equation approximation (3.11), but with a noise term added,

$$\frac{d\hat{x}}{dn} = n\,\delta\mu + a\hat{x}^2 + A\hat{\epsilon}(n), \tag{3.20}$$

where $\hat{\epsilon}(n)$ is white noise,

$$\langle\hat{\epsilon}(n)\rangle = 0, \quad \langle\hat{\epsilon}(n+n')\hat{\epsilon}(n)\rangle = \delta(n'),$$

and we have redefined the origin of the time variable n so that the parameter $\hat{\mu}$ sweeps through zero at $n = 0$ (i.e., we replaced n by $n - |\hat{\mu}_s|/\delta\mu$). Because we are only concerned with scaling, and not with the exact solution of (3.20), a fairly crude analysis will be sufficient.

First we consider the solution of (3.20) with the noise term omitted, and the initial condition [see (3.16)]

$$\hat{x}(0) = (-0.7290...)\left(\delta\mu/a^2\right)^{1/3}.$$

We define a characteristic point of the orbit, $\hat{x}_{\rm nl}(n_{\rm nl})$, where $a\hat{x}_{\rm nl}^2 \approx n_{\rm nl}\,\delta\mu$. For $n <$

$n_{\rm nl}$, $n\,\delta\mu \le d\hat{x}/dn < 2n\,\delta\mu$, and we can approximate the noiseless orbit as

$$\hat{x}(n) \approx \hat{x}(0) + \alpha(n)(n^2\,\delta\mu), \tag{3.21}$$

where $\alpha(n)$ is a slowly varying function of n of order 1 ($1/2 \le \alpha(n) < 1$ for $n < n_{\rm nl}$). Setting $a\hat{x}^2 \approx n\,\delta\mu$, we find that $n_{\rm nl}$ is given by

$$n_{\rm nl} \sim (a\,\delta\mu)^{-1/3}, \tag{3.22}$$

corresponding to [c.f., Eq. (3.21)]

$$\hat{x}_{\rm nl} \sim (\delta\mu/a)^{1/3}.$$

For $n > n_{\rm nl}$ (i.e., $\hat{x}(n) > \hat{x}_{\rm nl}$), Eq. (3.20) can be approximated as $d\hat{x}/dn \approx a\hat{x}^2$. Starting at $\hat{x}(n) \sim \hat{x}_{\rm nl}$, integration of this equation leads to explosive growth of $\hat{x}$ to infinity in a time of order $(a\,\delta\mu)^{-1/3}$, which is of the same order as $n_{\rm nl}$. Thus, the relevant time scale is $(a\,\delta\mu)^{-1/3}$ [this agrees with Eq. (3.18) in Sec. 3.2.5].

Now consider the action of noise. For $n < n_{\rm nl}$, we neglect the nonlinear term $a\hat{x}^2$, so that (3.20) becomes $d\hat{x}/dn = n\,\delta\mu + A\hat{\epsilon}(n)$. The solution of this equation is the linear superposition of the solutions of $d\hat{x}_a/dn = n\,\delta\mu$ and $d\hat{x}_b/dn = A\hat{\epsilon}(n)$, or $\hat{x}(n) = \hat{x}_a(n) + \hat{x}_b(n)$; $\hat{x}_a(n)$ is given by $\hat{x}_a(n) = \hat{x}(0) + n^2\delta\mu/2$, and $\hat{x}_b(n)$ is a random walk. Thus, for $n < n_{\rm nl}$, there is diffusive spreading of the probability density of $\hat{x}$,

$$\Delta_{\rm diff}(n) \equiv \sqrt{\langle \hat{x}_b^2(n)\rangle} \sim n^{1/2}A. \tag{3.23}$$

This diffusive spreading can blur out the structure in Fig. 3.6. How large does the noise amplitude A have to be to do this? We can estimate A by noting that the periodic structure in Figs. 3.6(b,c) results from orbits that take different integer times to reach

$\hat{x} \sim \hat{x}_{\rm nl}$. Thus, for $n \approx n_{\rm nl}$ we define a scale $\Delta_{\rm nl}$ in $\hat{x}$ corresponding to the periodicity in $1/\delta\mu$ by [c.f., Eq. (3.21)]

$$\hat{x}_{\rm nl} \pm \Delta_{\rm nl} \approx \hat{x}(0) + (n_{\rm nl} \pm 1)^2 \delta\mu$$

which yields

$$\Delta_{\rm nl} \sim n_{\rm nl}\delta\mu. \tag{3.24}$$

If by the time $n \approx n_{\rm nl}$, the diffusive spread of the probability density of $\hat{x}$ becomes as large as $\Delta_{\rm nl}$, then the noise starts to wash out the periodic variations with $1/\delta\mu$. Setting $\Delta_{\rm diff}(n_{\rm nl})$ from (3.23) to be of the order of $\Delta_{\rm nl}$ from (3.24), we obtain $n_{\rm nl}^{1/2} A \sim n_{\rm nl}\delta\mu$, which with (3.22) yields

$$A \sim (\delta\mu)^{5/6}. \tag{3.25}$$

Thus, we expect a collapse of the two parameter $(A, \delta\mu)$ data in Fig. 3.14(a) by means of a rescaling of A by $\delta\mu$ raised to an exponent 5/6 [i.e., $A/(\delta\mu)^{5/6}$].

3.3 Scaling of Indeterminate Saddle-Node Bifurcations for a Periodically Forced Second Order Ordinary Differential Equation

In this section we demonstrate the scaling properties of sweeping through an indeterminate saddle-node bifurcation in the case of the periodically forced Duffing oscillator [30],

$$\ddot{x} - 0.15\,\dot{x} - x + x^3 = \mu \cos t. \tag{3.26}$$

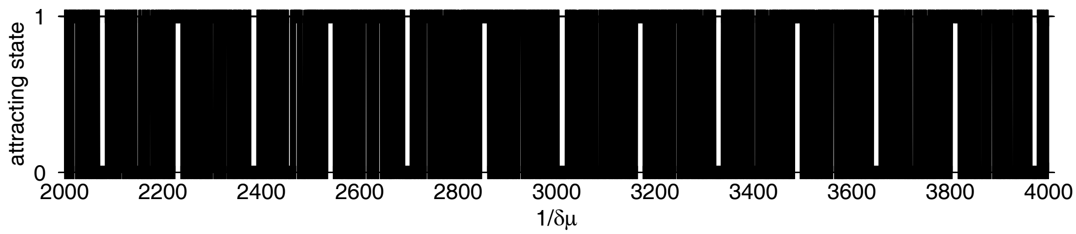

Figure 3.15: Final attracting state of swept orbits of the Duffing oscilator versus $1/\delta\mu$. The structure of white and black bands becomes asymptotically periodic. We have chosen $\mu_s = 0.253$, and $\mu_f = 0.22$. The attractor in the potential well for $x > 0$ is represented as a 1, and the attractor in the potential well for $x < 0$ is represented as a 0.

The unforced Duffing system (i.e., $\mu = 0$) is an example of an oscillator in a double well potential. It has two coexisting fixed point attractors corresponding to the two minima of the potential energy. For small μ, the forced Duffing oscillator has two attracting periodic orbits with the period of the forcing (i.e., 2π), one in each well of the potential. At $\mu = \mu_* \approx 0.2446$, a new attracting periodic orbit of period 6π arises through a saddle-node bifurcation. In Ref. [39], it is argued numerically that for a certain range of $\mu > \mu_*$ the basin of attraction of the 6π periodic orbit and the basins of attraction of the 2π periodic orbits have the Wada property. Thus, as μ decreases through the critical value μ_*, the period 6π attractor is destroyed via a saddle-node bifurcation on the fractal boundary of the basins of the other two attractors. This is an example of an indeterminate saddle-node bifurcation of the Duffing system which we study by considering the two-dimensional map in the $(\dot{x}, x)$ plane resulting from a Poincaré section at constant phase of the forcing signal. We consider orbits starting in the vicinity of the period three fixed point attractor, and, as we integrate the Duffing system, we decrease μ from $\mu_s > \mu_*$ to $\mu_f < \mu_*$ at a small rate of $\delta\mu$ per one pe-

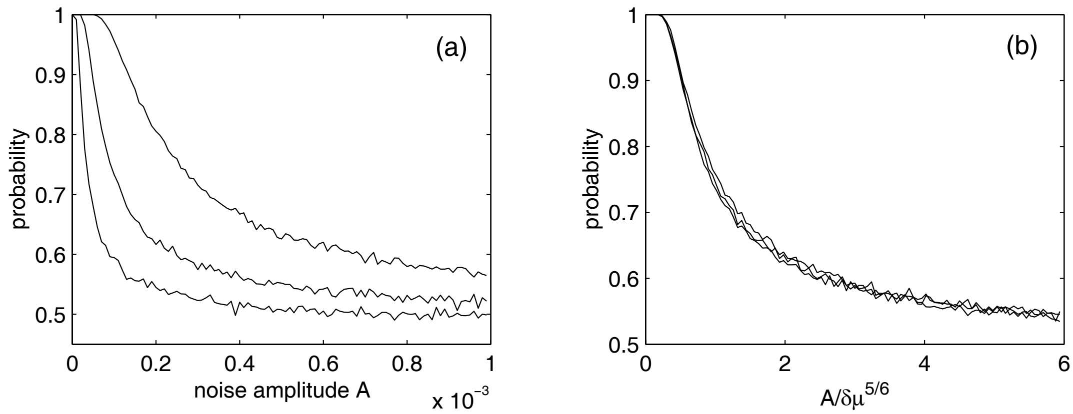

Figure 3.16: Probability the Duffing oscillator reaches the attracting periodic orbit in the potential well at $x > 0$ for three values of $\delta\mu$ spread over one decades: (a) versus the noise amplitude A, and (b) versus $A/(\delta\mu)^{5/6}$. We have chosen $\mu_s = 0.253$.

riod of the forcing signal. As μ approaches μ_*, (with $\mu > \mu_*$,) the period three fixed point attractor of the unswept Duffing system approaches its basin boundary, and the slowly swept orbit closely follows its location. For $\mu - \mu_* < 0$ small, the orbit will approximately follow the one-dimensional unstable manifold of the $\mu = \mu_*$ period three saddle-node pair. Thus, we can describe the sweeping through the indeterminate bifurcation of the Duffing oscillator by the theory we developed for one dimensional discrete maps. Figure 3.15 shows the final destination graph of a swept orbit initially situated in the vicinity of the period three fixed point of the Poincaré map. The final attracting state is represented as a 1 if situated in the potential well where $x > 0$, and is represented as a 0 if situated in the potential well where $x < 0$. As expected, the structure in Fig. 3.15 appears asymptotically periodic if graphed versus $1/\delta\mu$. In addition to slowly sweeping the Duffing system, consider an additive noise term $A\,\epsilon(t)$ in the right hand side of (3.26), where on every time step $\epsilon(t)$ is chosen randomly in $[-1, 1]$, and the time step used is $\Delta t = 2\pi/500$. Figure 3.16(a) shows the dependence of the probability of approaching the attractor represented as a 1 versus the noise amplitude A for three specially selected values of $\delta\mu$ (centers of white bands in the structure of Fig. 3.15 where the swept orbit reaches the attracting state represented by 1) spread over one decade. From right to left, the $\delta\mu$ values corresponding to the curves are approximately: 4.628716$\times 10^{-5}$, 1.461574$\times 10^{-5}$ and 4.621737$\times 10^{-6}$. Figure 3.16(b) shows collapse of the data in Fig. 3.16(a) to a single curve when the noise amplitude A is rescaled by $(\delta\mu)^{5/6}$, as predicted by our previous one-dimensional analysis (Sec. 3.2.6). Thus, we believe that the scaling properties of the indeterminate saddle-node bifurcation we found in one-dimensional discrete maps are also shared by higher dimensional flows.

3.4 Indeterminacy in How an Attractor is Approached

In this section we consider the case of a one dimensional map f_μ having two attractors A and B, one of which (i.e., A) exists for all $\mu \in [\mu_s, \mu_f]$. The other (i.e., B) is a node which is destroyed by a saddle-node bifurcation on the boundary between the basins of A and B, as μ increases through μ_* ($\mu_* \in [\mu_s, \mu_f]$). When an orbit is initially on B, and μ is slowly increased through μ_*, the orbit will always go to A (which is the only attractor for $\mu > \mu_*$). However, it is possible to distinguish between two (or more) different ways of approaching A. [In particular, we are interested in ways of approach that can be distinguished in a coordinate-free (i.e., invariant) manner.] As we show in this section, the way in which A is approached can be indeterminate. In this case, the indeterminacy is connected with the existence of an invariant nonattracting Cantor set embedded in the basin of A for $\mu > \mu_*$.

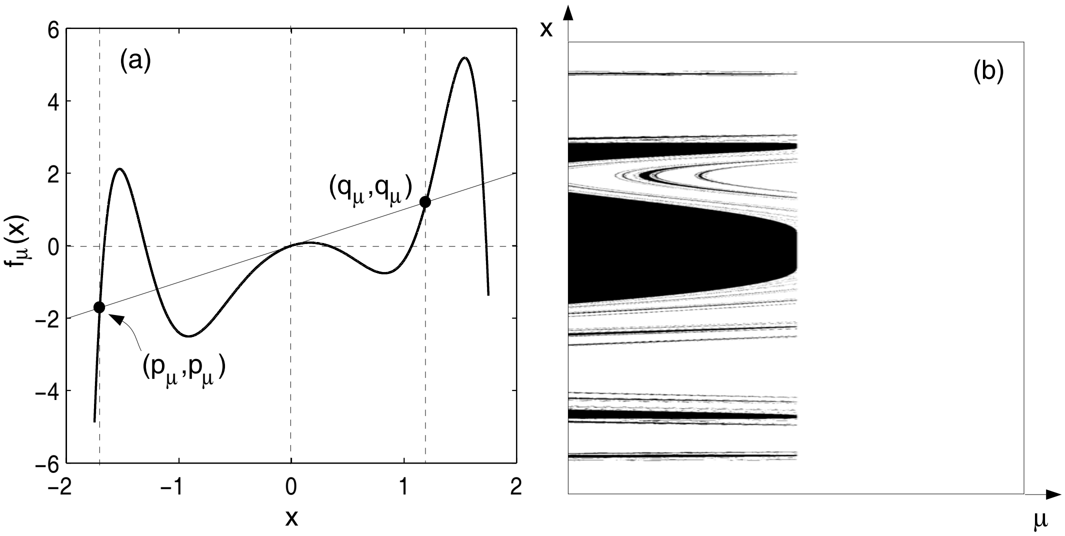

Figure 3.17: (a) Graph of $f_\mu(x)$ versus x at the bifurcation parameter. (b) Basin structure of map $f_\mu(x)$ versus the parameter μ ($-0.3 \leq \mu \leq 0.3$ and $-2 \leq x \leq 2$). The basin of attraction of the stable fixed point created by the saddle-node bifurcation is black while the basin of attraction of minus infinity is left white.

As an illustration, we construct the following model

$$f_\mu(x) = -\mu + x - 3x^2 - x^4 + 3.6x^6 - x^8. \tag{3.27}$$

Calculations show that f_μ satisfies all the requirements of the saddle-node bifurcation theorem for undergoing a backward saddle-node bifurcation at $x_* = 0$ and $\mu_* = 0$. Figure 3.17(a) shows the graph of f_μ versus x at $\mu = \mu_*$. Figure 3.17(b) shows how the basin structure of the map f_μ varies with the parameter μ. For positive values of μ, f_μ has only one attractor which is at minus infinity. The basin of this attractor is the whole real axis. As μ decreases through $\mu_* = 0$, a new fixed point attractor is created at $x_* = 0$. The basin of attraction of this fixed point has infinitely many disjoint intervals displaying fractal features [indicated in black in Fig. 3.17(b)]. This is similar to the blue basin $B[\mu]$ of the attractor B_μ of the previous one-dimensional

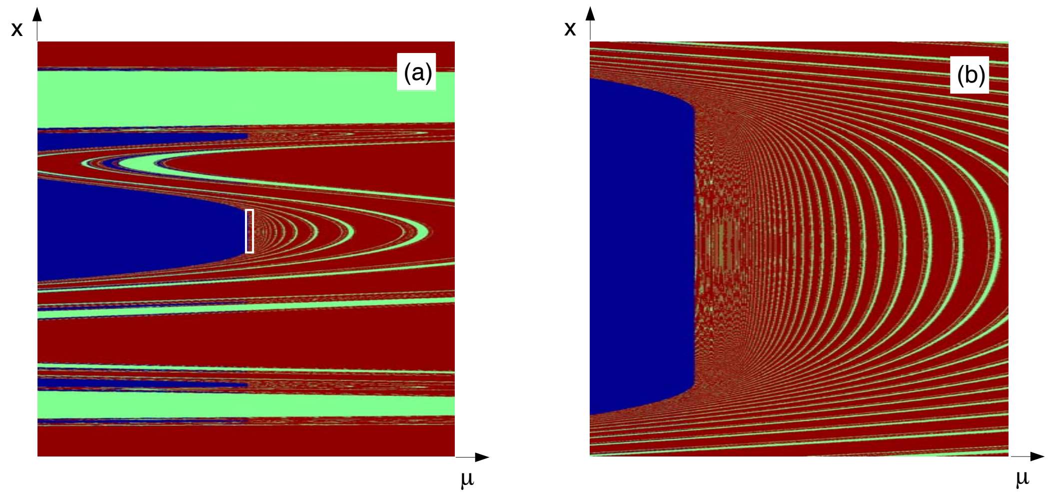

Figure 3.18: (a) Basin structure of f_μ versus μ ($-0.3 \leq \mu \leq 0.3$ and $-2 \leq x \leq 2$). We split the basin of attraction of minus infinity into two components, one plotted as the green region and the other plotted as the red region. The green region is the collection of all points that go to minus infinity and have at least one iterate bigger that the unstable fixed point q_μ. The red set is the region of all the other points that go to minus infinity. (b) Detail of Fig. 3.15(a) in the region shown as the white rectangle, $-0.005 \leq \mu \leq 0.015$ and $-0.09 \leq x \leq 0.41$.

model (see Sec. 3.2.1).

The blue region in Fig. 3.18(a) is the basin of attraction of the stable fixed point destroyed as μ increases through μ_*. For every value of μ we consider, the map f_μ has invariant Cantor sets. The trajectories of points which are located on an invariant Cantor set, do not diverge to infinity. One way to display such Cantor sets, is to select uniquely defined intervals whose end points are on the Cantor set. For example, Fig. 3.18(a) shows green and red regions. For every fixed parameter value μ, the collection of points that are boundary points of the red and green regions, constitutes

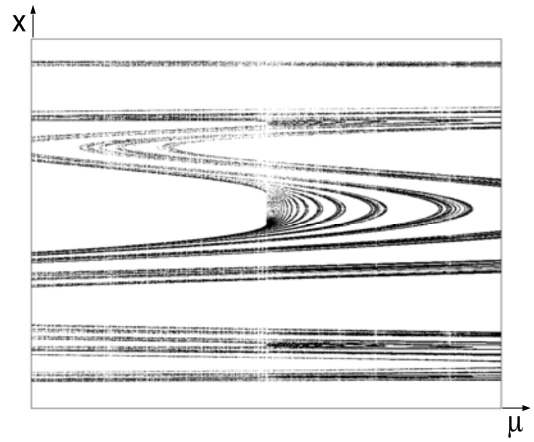

Figure 3.19: The chaotic saddle of f_μ versus μ ($-0.3 \leq \mu \leq 0.3$ and $-2 \leq x \leq 2$) generated by the PIM-triple method.

an invariant Cantor set. In order to describe these green and red regions, we introduce the following notations. For each parameter value μ, let p_μ be the leftmost fixed point of f_μ [see Fig. 3.17(a)]. For every $x_0 < p_\mu$, the sequence of iterates $\{x_n = f_\mu^{[n]}(x_0)\}$ is decreasing and diverges to minus infinity. For each value of μ, let q_μ be the fixed point of f_μ to the right of $x = 0$ at which $\frac{\partial f_\mu}{\partial x}(q_\mu) > 1$. A point $(x;\mu)$ is colored green if its trajectory diverges to minus infinity and it passes through the interval (q_μ, ∞), and it is colored red if its trajectory diverges to minus infinity and it does not pass through the interval (q_μ, ∞). Denote the collection of points $(x;\mu)$ that are colored green by $G[\mu]$, and the collection of points $(x;\mu)$ that are colored red by $R[\mu]$. Using the methods and techniques of [40], it can be shown that the collection of points $(x;\mu)$ which are common boundary points of $G[\mu]$ and $R[\mu]$ is a Cantor set $C[\mu]$ [41]. In particular, the results of [40] imply that for $\mu = \mu_* = 0$ the point $x_* = 0$ belongs to the invariant Cantor set $C[\mu_*]$.

Figure 3.18(b) is a zoom of Fig. 3.18(a) in the region of the saddle-node bifurca-

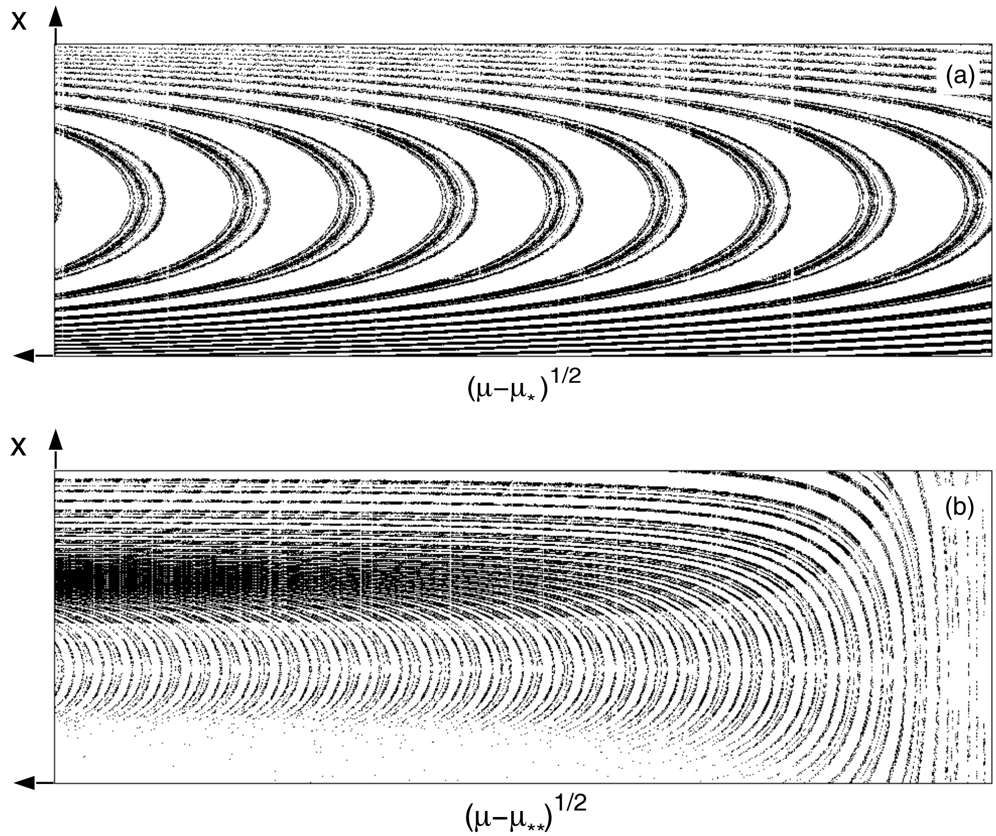

Figure 3.20: The chaotic saddle of the map f_μ in the vicinity of the saddle-node bifurcation with the horizontal axis rescaled from μ to: (a) $(\mu_* - \mu)^{-1/2}$. Notice that the chaotic saddle becomes asymptotically periodic ($-0.008 \leq x \leq 0.337$, $10 \leq (\mu_* - \mu)^{-1/2} \leq 15$). (b) $(\mu_{**} - \mu)^{-1/2}$, where $\mu_{**} = 0.23495384$. We believe that μ_{**} corresponds to the approximate value of the parameter μ where a saddle-node bifurcation of a periodic orbit of f_μ takes place on the Cantor set $C[\mu]$. In this case, the chaotic saddle also becomes asymptotically periodic ($-0.162 \leq x \leq 0.168$, $9.97 < (\mu_{**} - \mu)^{-1/2} < 2010$).

tion. For values of $\mu > \mu_*$, in the vicinity of $(x_*; \mu_*)$, one notices a fractal alternation of red and green stripes. The green and red stripe structure in Fig. 3.18(b) shares qualitative properties with the structure in Fig. 3.2(b). All the analysis in Sec. 3.2 can be adapted straightfowardly to fit this situation.

Figure 3.19 shows how the chaotic saddle of the map f_μ varies with μ. The chaotic saddle is generated numerically using the PIM-triple method. For an explanation of this method see Nusse and Yorke [42]. Using arguments similar to those in Sec. 3.2.3, we predict that changing the horizontal axis of Fig. 3.19 from μ to $(\mu - \mu_*)^{-1/2}$ makes the chaotic saddle asymptotically periodic. Numerical results confirming this are presented in Fig. 3.20(a). For f_μ given by (3.27), we were able to find a parameter value $\mu_{**} = 0.23495384$ where changing the horizontal axis of Fig. 3.19 from μ to $(\mu - \mu_{**})^{-1/2}$ [see Fig. 3.20(b)] apparently makes the chaotic saddle asymptotically periodic [with a different period than that of Fig. 3.20(a)]. As in the case discussed in Sec. 3.2, past the saddle-node bifurcation of f_μ at μ_*, infinitely many other saddle-node bifurcations of periodic orbits take place on the invariant Cantor set $C[\mu]$. We believe that μ_{**} is an approximate value of μ where such a saddle-node of a periodic orbit takes place.

3.5 Discussion and Conclusions

In this chapter, we have investigated scaling properties of saddle-node bifurcations that occur on fractal basin boundaries. Such situations are known to be indeterminate in the sense that it is difficult to predict the eventual fate of an orbit that tracks the pre-bifurcation node attractor as the system parameter is swept through the bifurcation. We have first analyzed the case of one-dimensional discrete maps. Using the normal form of the saddle-node bifurcation and general properties of fractal basin boundaries, we established the following universal (i.e., model independent) scaling results

- scaling of the fractal basin boundary of the static (i.e., unswept) system near the saddle-node bifurcation,
- the scaling dependence of the orbit's final destination with the inverse of the sweeping rate,
- the dependence of the time it takes for an attractor to capture a swept orbit with the -1/3 power of the sweeping rate,
- scaling of the effect of noise on the final attractor capture probability with the 5/6 power of the sweeping rate.

All these results were demonstrated numerically for a one-dimensional map example. Following our one-dimensional investigations, we have explained and demonstrated numerically that these new results also apply to two-dimensional maps. Our numerical example was a two-dimensional map that results from a Poincaré section of the forced Duffing oscillator. In the last section of this chapter, we have discussed how the new results listed above apply to the case where a saddle-node bifurcation occurs on an invariant Cantor set which is embedded in a basin of attraction, and we have supported our discussion by numerics.

BIBLIOGRAPHY

[1] A. S. Pikovsky, Sov. J. Commun Technol. Electron. **30**, 85 (1985).

[2] E. F. Stone, Phys Lett. A **163**, 367 (1992).

[3] A. Pikovsky, G. Osipov, M. Rosenblum, M. Zaks and J. Kurths, Phys. Rev. Lett. **79**, 47 (1997).

[4] E. Rosa, E. Ott and M. H. Hess, Phys. Rev. Lett. **80**, 1642 (1998).

[5] M. A. Zaks, E.-H. Park, M. G. Rosenblum and J. Kurths, Phys. Rev. Lett. **82** 4228 (1999).

[6] C. M. Ticos, E. Rosa, W. B. Pardo, J. A. Walkenstein and M. Monti, Phys. Rev. Lett. **85**, 2929 (2000).

[7] E. Allaria, F. T. Arecchi, A. DiGarbo and R. Meucci, Phys. Rev. Lett. **86**, 791 (2001).

[8] M. G. Rosenblum, A. S. Pikovsky and J. Kurths, Phys. Rev. Lett. **76**(11), 1804 (1996).

[9] P. Tass, M. G. Rosenblum, J. Weule, J. Kurths, A. Pikovsky, J. Volkmann, A. Schnitzler and H. -J. Freund, Phys. Rev. Lett. **81** 3291 (1998).

[10] F. Mormann, K. Lehnertz, P. David and C. E. Elger, Physica D **144** 358-369 (2000).

[11] A. Stefanovska, H. Haken, P. V. E. McClintock, M. Hožič, F. Bajrović and R. Ribarič, Phys. Rev. Lett. **85**(22), 4831 (2000).

[12] D. J. DeShazer, R. Breban, E. Ott and R. Roy, Phys. Rev. Lett. **87**(4), 044101 (2001).

[13] D. Maza, A. Vallone, H. Mancini and S. Boccaletti, Phys. Rev. Lett. **85**(22), 5567 (2000).

[14] O. Roessler, Phys. Lett. **57A**, 397 (1976).

[15] The parameter σ governs the amount of enhanced dispersion of the orbit circulation times. The chosen value $\sigma = 0.002$ is the same as that used in Ref. [4], and as can be seen from Fig. 2.2 of that reference, the added dispersion significantly effects the synchronization tongue in the vicinity of the chosen value of A $(A = 0.06)$.

[16] The length of the time series used to build the histogram graphs is four times larger than the length of the time series shown in Figs. 2.3(a1) and 2.3(a2), respectively. All histograms in this work have the same number of sample points distributed over 100 bins which makes an average of 2850 points/bin, ensuring good statistics for our purpose. (The expected percentage error for a typical bin is of the order of $100/\sqrt{2850} \approx 2\%$.)

[17] In this regard, we note that Fig. 2.1(a) only shows the region of locking parameters (A_0, T) in the limit of $t \to \infty$. Transients are not reflected by this figure. Thus, one should not expect that the slips occur exactely when $|\hat{A}|$ drops below

A_{th}. Futhermore, the tongue of *significant* weak synchronization is larger than the one for perfect synchronization displayed in Fig. 2.1(a).

[18] When $A_1 \gg A_2$, $\phi_H(t) \sim \phi_1(t)$ and when $A_2 \ll A_1$, $\phi_H(t) \sim \phi_2(t)$.

[19] S. W. McDonald, C. Grebogi, E. Ott and J.A. Yorke, Physica D **17**, 125 (1985).

[20] E. Ott, "Chaos in Dynamical Systems", second edition (Cambridge University Press, 2002).

[21] L.N. Virgin, "Introduction to Experimental Nonlinear Dynamics" (Cambridge University Press, 2000).

[22] J.M.T. Thompson and M.S. Soliman, Proc. R. Soc. Lond. A **432**, 101 (1991); M.S. Soliman and J.M.T. Thompson, Proc. R. Soc. Lond. A **438**, 511 (1992); J.M.T. Thompson, Physica D **58**, 260 (1992); M.S. Soliman, Proc. R. Soc. Lond. A **451**, 471 (1995).

[23] C. Baesens, Physica D **53**, 319 (1991); Lecture Notes in Mathematics **1493** (Springer-Verlag, Berlin, 1991), Ed. E. Benoît, 105; H.G. Davies and K. Rangavalhula, Proc. R. Soc. Lond. A **453**, 2043 (1997).

[24] N. Berglund and B. Gentz, Probab. Theory Relat. Fields **122**, 341 (2002).

[25] N.R. Lebovitz and A.I. Pesci, SIAM J. Appl. Math. **55**, 1117 (1995).

[26] J. Guckenheimer, Inventiones Math. **39**, 165 (1977); D. Whitley, Bull. London Math. Soc. **15**, 177 (1983); Y.A. Kuznetsov "Elements of Applied Bifurcation Theory", Applied Mathematical Sciences **112**, (Springer-Verlag 1995).

[27] Y. Pomeau and P. Manneville, Commun. Math. Phys. **74**(2), 189-197 (1980).

[28] J.-H. Cho, M.-S. Ko, Y.-J. Park and C.-M. Kim, Phys. Rev. E **65**, 036222 (2002); C.-M. Kim, O.J. Kwon, E.K. Lee and H. Lee, Phys. Rev. Lett. **73**, 525 (1994).

[29] H.E. Nusse, E. Ott and J.A. Yorke, Phys. Rev. Lett. **75**(13), 2482 (1995).

[30] R. Breban, H.E. Nusse and E. Ott, submitted for publication.

[31] J. Kennedy and J.A. Yorke, Physica D **51**, 213 (1991); H.E. Nusse and J.A. Yorke, Physica D **90**, 242 (1996); H.E. Nusse and J.A. Yorke, Science **271**, 1376 (1996); Ergodic Th. & Dynam. Syst. **17**, 463 (1997).

[32] C. Robert, K.T. Alligood, E. Ott and J.A. Yorke, Physica D **144**, 44 (2000).

[33] B.-S. Park, C. Grebogi, E. Ott and J.A. Yorke, Phys. Rev. A **40**(3), 1576 (1989).

[34] M. Misiurewicz and A.L. Kawczynski, Commun. Math. Phys. **131**, 605 (1990).

[35] H.E. Nusse and J.A. Yorke, Commun. Math. Physics **150**, 1 (1992).

[36] J. Palis and M. Viana, Lecture Notes in Mathematics **1331** (Springer-Verlag, Berlin, 1988), Eds. R. Bamón, R. Labarca and J. Palis, 150; F. Takens, ibid., 196.

[37] "Handbook of Mathematical Functions with Formulas, Graphs, and Mathematical Tables", Eds. Milton Abramowitz and Irene A. Stegun, Supt. of Docs., U.S. Govt. Print. Off., Washington, 1970.

[38] For numerical purposes, we consider that a swept orbit has reached its final state if $|x_{n+1} - x_n| < 10\,\delta\mu$. Changing the factor of 10 to other values ranging from 1 to 30, can slightly change the numerical result for the slope (by as much as 3%).

[39] J. Aguirre and M.A.F. Sanjuán, Physica D **171**, 41 (2002).

[40] H.E. Nusse, SIAM J. Appl. Math. **47**, 498 (1987).

[41] For every μ ($-0.3 < \mu < 0.3$), write $p^*_\mu = \max\{x \in \mathbb{R} : f_\mu(x) = p_\mu\}$ and $q^*_\mu = \max\{x \in \mathbb{R} : f_\mu(x) = q_\mu\}$. The interval $[q_\mu, q^*_\mu]$ also contains a Cantor set. By coloring this whole segment green, this information is lost. Therefore, the coloring scheme should be adapted if one wants to have the whole invariant Cantor set represented for every μ. For example, if a trajectory that diverges to minus infinity contains a point that is greater than p^*_μ then the initial point is colored green, if a trajectory that diverges to minus infinity contains a point that is greater than q^*_μ but not greater than p^*_μ then the initial point is colored yellow. A point is colored red, if its trajectory diverges to minus infinity and does not have a point that is greater than q^*_μ. Then the collection of boundary points (a point x is a boundary point if every open neighborhood of x contains points of at least two different colors) is a Cantor set that contains the Cantor set described above.

[42] H.E. Nusse and J.A. Yorke, Physica D **36**, 137 (1989).