Part I

Basic Concepts of Nonlinear Dynamics

2

An Overview of Nonlinear Phenomena

In Part I we aim to give a general outline of nonlinear dynamics, which is an essential prerequisite to our more advanced studies including our goal of understanding chaotic motions. This chapter provides a quick overview of the nonlinear dynamics field, before we begin our more detailed presentation.

2.1 UNDAMPED, UNFORCED LINEAR OSCILLATOR

We start our overview by looking at the undamped, unforced linear oscillator of Figure 2.1. The equation chosen for this first illustration has the stiffness constant $4\pi^2$, which makes the periodic time equal to unity. The solution of such an equation is simply a sine wave, the constant amplitude and phase of which are determined by the starting values of x and $\dot{x}$. So, once started, we have a constant sine wave that persists for all time, and there is no transient or

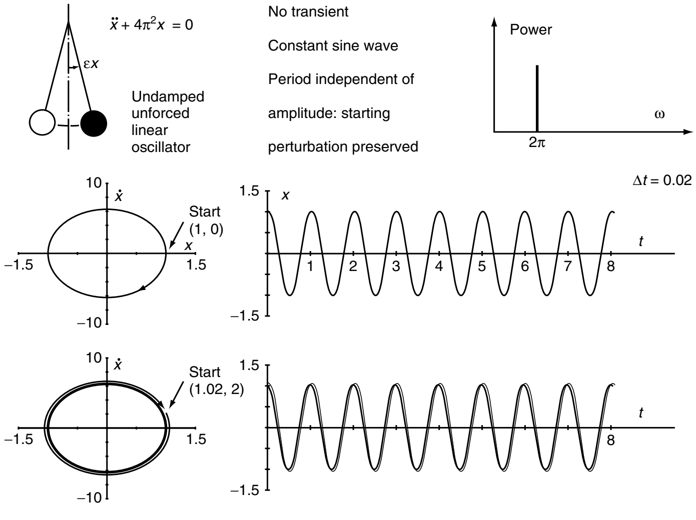

Figure 2.1 Undamped, unforced behaviour of a linear oscillator

decay of any kind. The periodic time, unity in the present example, is a constant independent of the starting conditions, the amplitude of the motion, and the time.

A typical plot of x against the time t is shown, resulting from the starting condition $(x, \dot{x})$ equal to (1, 0) at the time $t = 0$. Along with the other, mainly nonlinear, problems considered in this chapter, this solution was obtained by numerical time integration using a fourth-order Runge–Kutta routine on a desktop Hewlett-Packard computer with the step size indicated, here $\Delta t = 0.02$.

If we plot not x against t but $\dot{x}$ against x, we have the phase portrait shown on the left. Starting as before at (1, 0) we now have the closed ellipse shown, the representative point moving continuously round and round this closed orbit as the time goes to infinity. The power spectrum of this response, shown in the top right-hand diagram, is simply a spike (or delta function) at the circular frequency of 2π radians per second.

We must finally ask the question: what would happen if we changed the starting condition by a small amount? The answer is illustrated in the lower diagram, where we show both the *fundamental* reference motion starting at (1, 0) and a perturbed motion starting at (1.02, 2). We see that we have two sine waves running in step with just a small difference in amplitude and phase resulting from the slightly different starting values of x and $\dot{x}$. They continue to run nicely in step for all time because the period of oscillation of the two motions is the same (and equal to unity, as we have seen). So a starting perturbation is preserved, and the fundamental motion is *neutrally stable* in a dynamical sense.

In the left-hand phase space, the two motions appear as neatly nesting ellipses. All possible motions of this linear oscillator are indeed represented by a complete family of nesting ellipses, which represent the full phase portrait of the system. The orbit passing through any particular starting point $(x, \dot{x})$ defines the subsequent unique motion of the oscillator.

This linear oscillator models in an approximate fashion many basic physical systems, such as for example the free motions of a simple hanging pendulum. The modelling is however unrealistic in two important ways. First, it ignores the damping action of inevitable dissipative forces, such as air resistance in the example of the laboratory pendulum. In the absence of impressed driving forces, the motions of all real macroscopic mechanical systems will eventually decay, as with a free experimental pendulum, so our present equation fails to model this vital aspect. Secondly, all real systems will have some degree of nonlinearity, which in itself modifies the behaviour in important ways. Large-amplitude oscillations of an undamped pendulum are for example governed by a nonlinear differential equation that we shall examine next: a linear approximation to the behaviour of a pendulum is only valid for small angles of oscillation.

The two unrealistic approximations of *linearized stiffness* and *zero damping* will be removed in turn, so we look next at the large-amplitude, nonlinear motions of an undamped pendulum.

2.2 UNDAMPED, UNFORCED NONLINEAR OSCILLATOR

The undamped, unforced nonlinear system of Figure 2.2 represents the *exact* equation of motion of a simple pendulum undergoing arbitrarily large oscillations. This equation in terms of the angle x is easily derived using Newton's law of motion for the bob by resolving perpendicular to the light string to eliminate the unknown tension: alternatively it can be derived by Lagrangian or Hamiltonian energy methods. The length of the pendulum, relative to the gravitational constant, has been chosen to make the coefficient equal to $4\pi^2$. So for small oscillations we could *linearize* the equation by approximating sin x to x, and retrieve the linear oscillator of our earlier discussion, with periodic time equal to unity.

The solution of this nonlinear differential equation can be obtained after some algebra in terms of elliptic integrals: alternatively the equation can be easily integrated numerically on a digital computer as we have done here. Depending on the starting conditions of $(x, \dot{x})$ we now find a steady undamped oscillation corresponding to the motion of our idealized undamped pendulum. A *given motion* from a given start thus exhibits no transient or decay, just a steady waveform of constant amplitude and constant period. The waveform is not however sinusoidal, and could in fact be decomposed by Fourier analysis into a fundamental harmonic plus odd higher harmonics: this gives rise to the power spectrum shown with a large spike at a certain circular frequency ω_F and smaller spikes at 3, 5, 7, . . . times this value.

The central waveform shows the steady oscillation starting at (3.054, 0) corresponding to the pendulum starting from rest with $\dot{x} = 0$ at a value of $x = 3.054 \times 180/\pi = 175^\circ$. To visualize this physically we must suppose that the heavy pendulum bob is supported not by a string, which could become slack, but by a light rigid rod pivoted to the fixed support. Because this

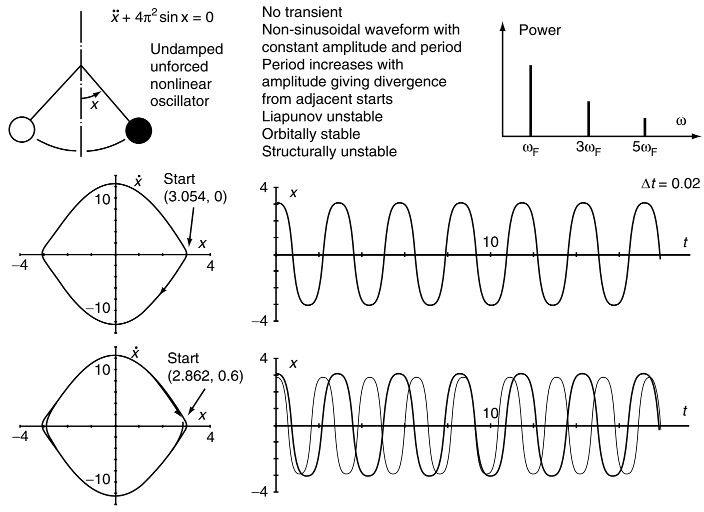

Figure 2.2 Undamped, unforced behaviour of a nonlinear oscillator

rigid-link pendulum would be in (unstable) equilibrium at $x = 180^\circ$, the motion begins very slowly and the waveform is very flat and quite noticeably non-sinusoidal. The corresponding $\dot{x}(x)$ phase picture is shown to the left-hand side: the closed trajectory is quite clearly not elliptical, and has a high curvature on the x axis corresponding to the proximity of an unstable equilibrium state.

Now the periodic time of a given motion is constant as we have just seen, but the period of different motions *increases* with the amplitude. It is clear for example that a start very close to $x = 180^\circ$ will give a motion with a very large period, since at the end of each big swing the pendulum will almost come to rest in the inverted position: indeed the periodic time goes to infinity as the amplitude approaches π. Notice that the periodic time of our displayed waveform is about 3, compared with the periodic time of unity for the small-amplitude linearized motions.

This variation of period with amplitude gives rise to a new phenomenon when we consider a perturbed motion. The lower diagram shows the fundamental motion just considered together with a perturbed motion starting from slightly different initial conditions. Because these new conditions give rise to a motion with a slightly different amplitude, the perturbed waveform has a slightly different period. So we have a *beat* phenomenon and the two motions drift in and out of phase with one another. This means that, although the two waveforms will eventually resynchronize, there is an initial *divergence* from adjacent starts. This makes the fundamental oscillatory motion unstable in the strict sense of Liapunov. In the left-hand phase diagram however, in which the *time* discrepancies of the two motions are not visible, the two closed *orbits* are seen to lie everywhere close to one another: in recognition of this fact the fundamental motion is said to be *orbitally stable*.

For the motions under consideration, the phase portrait of the present undamped nonlinear oscillator consists of nesting closed orbits. For small oscillations these are roughly elliptical corresponding to the nearly sinusoidal waveform, but they become increasingly distorted with increasing curvature near the x axis for the larger non-sinusoidal motions.

The steady undamped oscillations of our first two examples are not typical of real undriven systems. Clearly the smallest trace of dissipation will give damped waveforms, and the nest of closed orbits in the phase space will become *inward spirals*. The fact that the topological nature (closure) of the phase orbits can be destroyed by even infinitesimal damping is recognized by declaring the pathological undamped systems to be *structurally unstable*.

For the rest of this chapter we shall be concerned with *typical* damped systems, and we start by looking at the behaviour of a damped linear system.

2.3 DAMPED, UNFORCED LINEAR OSCILLATOR

We consider then the differential equation of Figure 2.3, which is written in a rather standard form, with ζ representing the damping factor, namely the ratio of the actual damping to the critical damping at which oscillatory behaviour

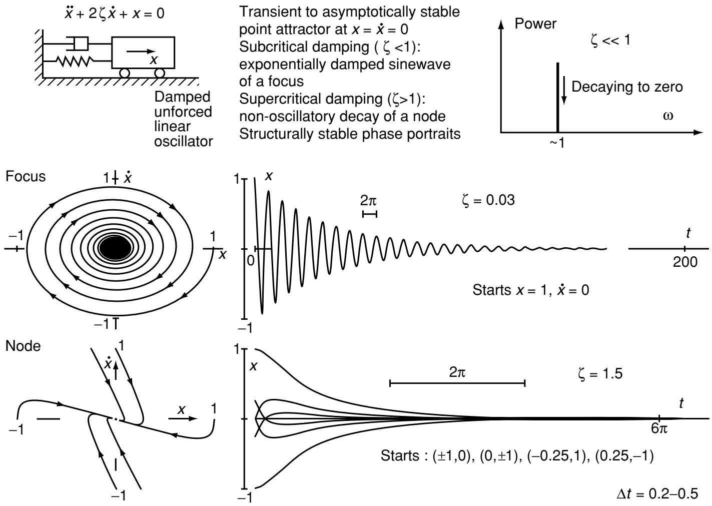

Figure 2.3 Damped, unforced behaviour of a linear oscillator

ceases. We can think of this equation as representing the motion of a mass constrained by a linear elastic spring in parallel with a dashpot full of oil, which is assumed to provide a force opposing the instantaneous velocity.

An analytical solution of this linear differential equation is readily written down: for light damping with $\zeta < 1$ we have an exponentially damped sine wave, while for heavy damping with $\zeta > 1$ we have a non-oscillatory exponential decay.

A typical lightly damped waveform is shown in the middle picture, starting at $x = 1$, $\dot{x} = 0$. The decaying wave has a constant period, defined for example by successive crossings of the time axis, which is nevertheless slightly dependent on the value of ζ. With the light damping shown, the period is essentially unchanged from the period, 2π, of the corresponding undamped system obtained by setting $\zeta = 0$. For light damping the power spectrum will be roughly a single spike decaying to zero along with the wave amplitude.

The corresponding phase portrait on the left is now a spiral, heading inwards towards the asymptotically stable equilibrium state at the origin (0, 0). The full linear phase portrait, termed a focus, is a set of intertwining, non-crossing spirals. Every motion here represents a transient to the asymptotically stable equilibrium state of rest at the origin, which for obvious reasons is called a *point attractor*. The whole phase portrait is now *structurally stable* since for finite damping the spiralling form cannot be topologically changed by *any* infinitesimal changes to the system.

The pictures for heavier supercritical damping shown below give the waveform and phase trajectories for six alternative starts. The system moves back to its stable state of rest in a direct non-oscillatory fashion, and the whole phase portrait is called a *node*. Once again, we have a structurally stable point attractor at (0, 0) capturing all motions of the system.

Since all motions decay to rest, fundamental and perturbed motions coalesce as time goes to infinity, and starting perturbations are lost.

2.4 DAMPED, UNFORCED NONLINEAR OSCILLATOR

To conclude our examination of unforced (undriven) systems, we look now at a damped nonlinear problem, typified by the pendulum of Figure 2.4. This is the large-amplitude pendulum of our earlier discussion, now with the modelling of air drag by a realistic velocity-squared law: notice that the damping force proportional to $\dot{x}^2$ has to be entered into the differential equation of motion as $\dot{x}|\dot{x}|$ to ensure that it is always opposing the velocity. Having put on this quadratic damping, we should perhaps emphasize that the form of damping is largely irrelevant to the following discussion, the salient points being just as well illustrated by the use of linear damping: the computed traces relate however to the quadratic damping.

Clearly we once again have transients to the asymptotically stable hanging equilibrium state representing a point attractor in the phase space.

The central waveform damps and becomes increasingly sinusoidal as x becomes small, while the power spectrum is a decaying set of spikes as shown. The phase portrait is a spiral, becoming increasingly elliptical as the trajectories approach the central attractor. A little linear damping would be needed to make this portrait structurally stable near the origin.

As with the undamped pendulum suffering large-amplitude oscillations, adjacent starts still exhibit a temporary beating character with an associated initial divergence due to the variation of the period with amplitude. But initial perturbations are eventually lost as all motions coalesce in the unique hanging state.

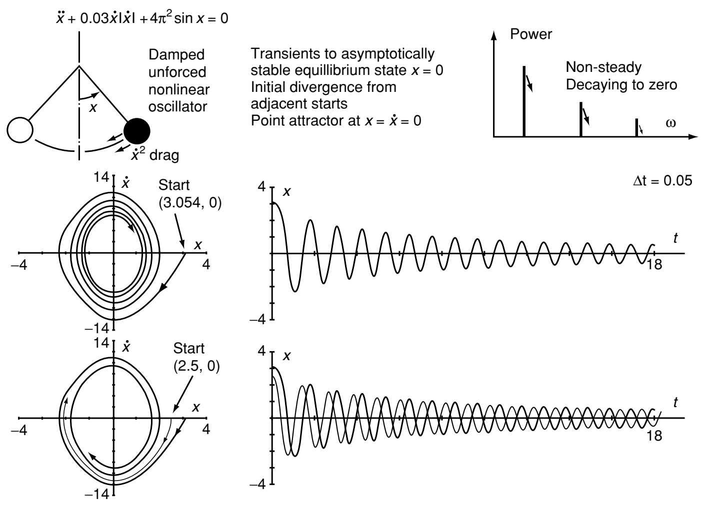

Figure 2.4 Damped, unforced behaviour of a nonlinear oscillator

This local phase portrait is a set of intertwining spirals with all motions captured by the central attractor. The full phase portrait of a pendulum including high-velocity motions passing through the inverted state is most nicely seen in a cylindrical phase space, and will be presented later (Figure 3.15).

2.5 FORCED LINEAR OSCILLATOR

We have so far looked only at autonomous unforced systems with zero on the right-hand side of the equation, but we turn now to sinusoidally driven non-autonomous oscillators. Damping, we have seen, is an essential ingredient of good modelling, so we shall start by looking at the damped, forced linear oscillator of Figure 2.5. This would be an adequate mathematical model of a pin-ended steel beam driven to small-amplitude lateral oscillations by an electromagnet carrying a sinusoidal alternating current. Here physical damping would arise from air resistance and internal material dissipation. The numerical coefficients have been chosen to provide a sharp frequency contrast between the transient and the steady-state solution, and the damping ratio of the unforced left-hand side is 0.1 (see equation 3.11).

This is a classical resonance problem of engineering texts, and the well-known analytical solution is easily written down. It is the algebraic sum of the so-called particular integral (PI) and the complementary function (CF). The CF is just the solution obtained by setting the left-hand side of the equation to zero: that is to say it is the exponentially damped sinusoidal solution of the unforced autonomous system. It has the usual two arbitrary constants of amplitude and phase obtained by applying the starting conditions to the *whole* solution. With the present choice of constants the CF is a high-frequency sine wave with quite a heavy rate of damping.

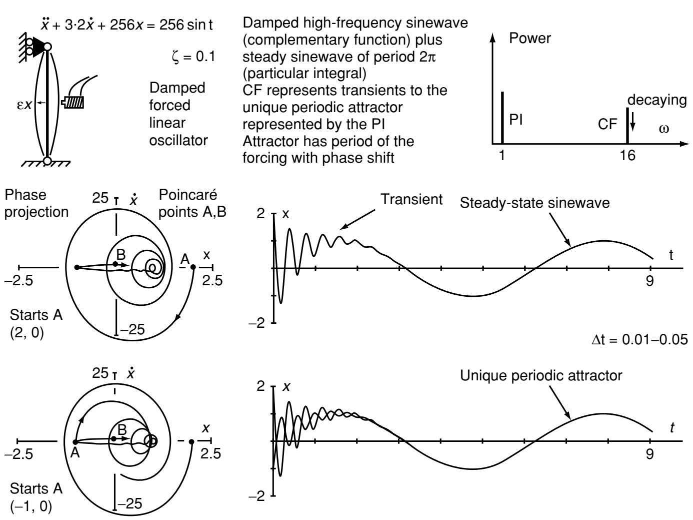

Figure 2.5 Damped, forced behaviour of a linear oscillator

The PI is a particular (known) solution of the whole equation, being in fact a steady undamped sine wave with the same frequency as the forcing term with which it has a fixed phase difference. The amplitude of the PI depends crucially on the ratio of the forcing frequency to the natural frequency of the autonomous left-hand side, being large when this ratio is close to one so that we have a condition of resonance. The conventional engineering resonance response curves simply plot the magnitude of the PI against this frequency ratio, giving for light damping a sharp peak at unity.

We should emphasize here, however, that from the qualitative dynamics point of view it is irrelevant whether the system is 'at resonance' or not. With the particular coefficients chosen, our illustration is well away from the resonant condition, but the discussion of the system's behaviour is essentially unrelated to this fact.

Since the analytical solution is just the algebraic sum of the CF and the PI, it is clear that the former damped sine wave represents a decaying transient, which leaves the PI as the unique final steady state: this is the reason for the engineer's consuming interest in the amplitude of the PI. A waveform starting at (2, 0) is shown in the central figure and we see clearly the high-frequency transient leading rapidly to the steady sinusoidal state described by the PI.

Now a forced system such as this has a three-dimensional phase space defined by the coordinates $(x, \dot{x}, t)$, the essence of phase spaces being that they are full of *non-crossing* trajectories. It is sometimes convenient, however, just to plot the *phase projection* $(x, \dot{x})$ and accept the fact that trajectories will appear to cross in this projection. The phase projection corresponding to the drawn waveform is thus shown to the left-hand side. The high-frequency transient appears as decaying circles, and the final steady state as a very long, thin ellipse pointing along the x axis.

It is also helpful in the phase projection to make a dot, or small circle, whenever the forcing cycle is about to commence, at t equal to multiples of the forcing period, here 2π. This is the so-called Poincaré section and is represented by points A and B in the present time integration. Since the final steady state is here an oscillation with the same period as the forcing, the final steady-state mapping will be the constant repetition of a fixed point, here quite close to B. *Mapping* from section to section is defined in Figure 5.1.

The lower pictures show, superimposed, the effect of a completely different start. As dictated by the analytical solution, the different transients resulting from different integration constants in the CF lead merely to the same *unique periodic attractor* corresponding to the PI. As we have seen, this attractor is sinusoidal with the period of the forcing, but with a constant phase shift. The power spectrum will be predominantly two spikes at the forcing frequency and at the natural autonomous frequency, the latter decaying as the transient is lost.

2.6 FORCED NONLINEAR OSCILLATOR: PERIODIC ATTRACTORS

Just as a stiffness nonlinearity introduced new phenomena into the response of an unforced oscillator, so a nonlinearity generates new features in a driven system. So we look now at the damped, forced nonlinear oscillator illustrated in Figure 2.6. This is the sinusoidally (here cosinusoidally) forced Duffing equation with a linear and a cubic stiffness. This could be used to model the moderately large bending deflections of an electromagnetically driven steel beam held pinned to fixed supports as shown. These fixed supports induce a membrane tension at finite deflections, which gives a hardening nonlinear stiffness modelled for moderately large deflections by the cubic term.

For such a driven nonlinear oscillator, closed-form analytical solutions are not available and recourse *must* inevitably be made to numerical time integrations. Just as with the preceding linear system, transients are observed, but after these have decayed we now find that there are two alternative stable steady states denoted here by A and B. The first plot of x against t shows these two steady oscillatory states, the starting points to eliminate transients having been found by previous trial computations. We see that the large-amplitude motion A and the small-amplitude motion B both have the same period as the forcing term and are therefore fundamental *harmonics* as opposed to subharmonics: they are noticeably out of phase with one another. The corresponding steady-state phase projections are shown in the left-hand phase diagram, each closed orbit having one Poincaré mapping denoted by a circle because the motions have the period of the forcing: these mapping points show where the system is whenever the time is a multiple of 2π.

These two steady-state solutions, A and B, can be seen on the resonance response diagram at the top right. This is a plot of the response amplitude against the ratio of the forcing frequency to the natural frequency of the autonomous system: this ratio is 1.6 for the parameters here adopted. Now in a linear resonance problem we have a vertical resonant peak, but the positive

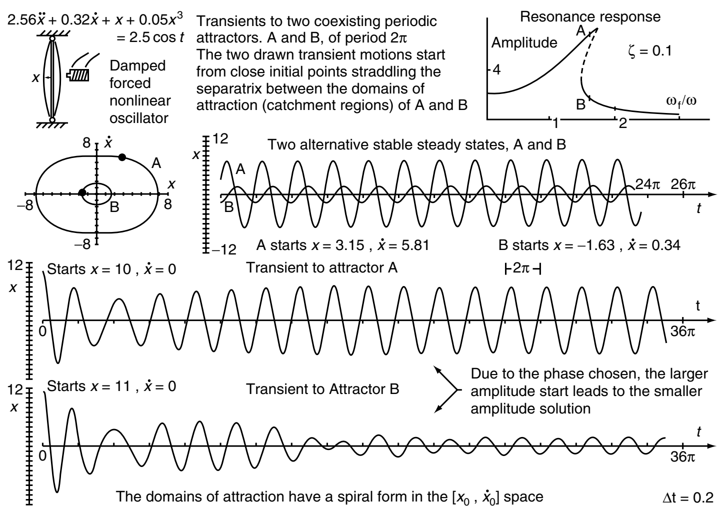

Figure 2.6 Damped, forced behaviour of a nonlinear oscillator: transients to periodic attractors

cubic stiffness of our Duffing's equation curves the peak to the right, giving a domain of frequency ratio with three steady states. The steady state of intermediate amplitude is unstable and so is not observed in a normal time integration, leaving us with the two alternative stable solutions A and B.

Now which of these two coexisting periodic attractors is picked up in a given time integration depends on the starting conditions, and two transient motions are illustrated in the lower diagrams. Starting with $(x, \dot{x})$ equal to (10, 0) gives a transient leading to attractor A while starting at (11, 0) gives a transient leading to attractor B. Notice that due to the phase chosen, the *larger*-amplitude start leads to the *smaller*-amplitude solution. The more obvious converse could equally apply, the final motion adopted being as much governed by phase as by amplitude.

Clearly in the space of the starting values of $(x, \dot{x})$ at $t = 0$ there will be *basins of attraction* such that motions originating in the basin of A lead after the decay of transients to solution A, while motions starting in the catchment region of B lead to the periodic attractor B. Between the basins of attraction (catchment regions) will be a separatrix curve, and it is clear that our two rather close starts straddle this separator. The basins of attraction tend to have a complex spiral form, which accounts for the sensitivity to both phase *and* amplitude previously mentioned.

This multiplicity of alternative stable attracting solutions (often more than the present two) dependent on the starting conditions, which is not encountered in the linear resonance problem with its unique periodic attractor, is typical of nonlinear driven oscillators.

We come at last to our final equation of this chapter giving rise, as the reader might expect, to a chaotic solution governed by a strange attractor.

2.7 FORCED NONLINEAR OSCILLATOR: CHAOTIC ATTRACTOR

The system of Figure 1.7, discussed briefly in Chapter 1, is a version of the driven Duffing equation studied extensively by Ueda, and we see that it differs from our previous damped, forced nonlinear oscillator in having no linear stiffness. This would in fact arise physically if we had a beam loaded to precisely its (Euler) buckling load: at buckling the linear stiffness has dropped to zero due to the destabilizing action of the axial compressive load, and the nonlinear stiffness can be modelled locally by the cubic term.

Once again analytical solutions are impossible, and digital computations show that after transients have decayed the system settles down to a condition of steady-state chaos. In contrast to the point and cyclic attractors that we have so far examined, this convergence to chaos is said to be governed by a *chaotic attractor*. These chaotic or strange attractors can coexist with other periodic steady states, with appropriate basins of attraction, etc., but for the coefficients chosen here there is in fact just a unique chaotic attractor that captures all motions of the system. The middle trace shows a rather brief but fairly obvious transient from (0, 0) lasting visibly for only about five forcing cycles of period 2π. The steady-state chaos covering the remaining 45 forcing cycles has a fairly regular though non-periodic appearance, and we notice that the positive x peaks synchronize approximately with the start of a forcing cycle for which t is a multiple of 2π.

The steadiness of this final chaotic state is reflected in a stationary power spectrum and a *typical* spectrum of chaos is shown in the upper right picture. This is due to Ueda, and is for a slightly different set of coefficients, with 0.1 and 12 replacing the 0.05 and 7.5 of our equation. We see spikes at the forcing frequency and odd multiples of this frequency (typical of a non-sinusoidal periodic wave with the period of the forcing) plus, however, regions of 'white noise' extended broadband peaks.

The bottom trace shows a more dramatic transient, generated by starting at large amplitude at an inconvenient phase. The high frequency is a natural consequence of the large x, because the effective stiffness increases as x^2. However, even after this start, the recognizable pattern of the steady-state chaos soon emerges.

We recollect that the phase space of this driven oscillator is three-dimensional, spanned by $(x, \dot{x}, t)$, and the Poincaré mapping is generated by the successive intersections of a trajectory with the $t = 2i\pi$ sections, where $i = 0, 1, 2, \ldots$. The steady-state chaotic mapping is shown in the last picture. Here the dots build up to form a complete shape with a fractal structure, similar to that of a Cantor set (Figure 11.9). All the points lie in the positive x regime, corresponding to our earlier observation that in the final state the positive x peaks synchronize with the beginning of the forcing cycle.

Transients would appear as rather scattered dots outside this attractor, but as we have seen the mapping points are very quickly attracted into this set. The Poincaré section, often itself referred to as the attractor of the chaotic motion, is really just a cross-section of the full attracting structure, which is a fixed geometric form in the full three-dimensional phase space to which all trajectories are finally attracted. It is the continuous stretching and folding of the sheets of this attractor that produces the turbulent mixing motions characteristic of chaotic dynamics, as we shall see in Chapter 6.