Safety criteria for aperiodically forced systems

Z.Bishnani
R.S.MacKay

Mathematics Institute,
University of Warwick, Coventry CV4 7AL, U.K.

16 Feb 2001; revised 10 Mar 2002

Abstract

The theory of uniform hyperbolicity is used to obtain sets of forcing functions for which the response of a dynamical system remains in a prescribed "safe" region. The resulting safety criteria are asymptotically optimal as the size of the safe region goes to zero.

AMS codes: 93C73, 34D10
keywords: aperiodic forcing, uniform hyperbolicity, safety criteria, non-autonomous attractors

1 Introduction

The theory of non-autonomous dynamical systems is highly relevant to disciplines like engineering, medicine and economics, in which the system of interest is often subject to time-varying inputs of a general form. For just one example of area of application, see [BA].

It has old roots, e.g. [MS, Se, Fr, SS, Cop, J], but there has been a large amount of recent activity, e.g. [MSe, RGO, CK1, CV, Gu, Arn, JN, CK, CKS, Sie, Sim, LRS], session $C2$ of [FGS] and references contained in these.

We present here some ideas (worked out mainly in 1993–5; more details appear in [Bi]), on how to use the theory of uniform hyperbolicity to determine rigorous "safety criteria" for aperiodically forced systems, i.e. conditions on forcing functions for which the response is guaranteed to remain within a specified "safe" region. If the safe region is a norm-ball, our method is asymptotically optimal as the radius tends to zero.

We achieve this in two situations. The first is where for no forcing, the system has a linearly attracting equilibrium in the prescribed region. The second is where for a particular "reference" periodic forcing, the system has a linearly attracting periodic solution to which it is desired to remain close (up to suitable time-reparametrisation). The former is relevant to a wide range of applications, ship capsize being a notable example. A primary motivation for the second is phase-locked loops: oscillators which are supposed to track the phase of an oscillatory input near the natural frequency but not necessarily exactly periodic.

The method can also be used to estimate explicit neighbourhoods of attraction for linearly attracting solutions, and to derive safety estimates for permissible deviations from design of a dynamical system, though these are not our main emphases here. It applies equally well to feedback control systems, by considering the controlled plus controlling system as one joint dynamical system subject to inputs. The forcing can be additive, parametric, or any combination.

The plan of the paper is as follows. In §2, we recall the relevant elements of uniform hyperbolicity theory and explain how it can be used to obtain rigorous safety criteria. In §3 we present a refinement of the method, which is asymptotically optimal for small safe regions. In §4 and §5, we apply the method to damped nonlinear oscillators, firstly subject to additive forcing and then allowing also parametric forcing. In §6, we apply the method to phase-locked loops. In §7 and §8 we outline extensions to determine neighbourhoods of attraction and to allow discontinuous or distributional forcing. The papers ends with concluding remarks in §9.

2 Uniform hyperbolicity

An equilibrium x_0 of an autonomous dynamical system $\dot{x} = f(x)$ on $\mathcal{R}^n$ (or more generally a Banach manifold M) is called *hyperbolic* if the derivative Df_{x_0} exists and has no spectrum on the imaginary axis. A special case is when all the spectrum is in the open left-half plane, called *linearly attracting.*

For a time-periodic dynamical system $\dot{x} = f(x,t) = f(x,t+T)$ of period $T > 0$, a periodic solution $x_0(t)$ of period kT, for some $k \in \mathcal{N}$, is called *hyperbolic* if the monodromy map M from $t = 0$ to $t = kT$ is defined and has no spectrum on the unit circle. It is called *linearly attracting* if all the spectrum is in the open unit disk.

To motivate a generalisation of hyperbolicity to general solutions of general ODEs, it is convenient to unify these definitions using the following two theorems. To simplify presentation we take $M = \mathcal{R}^n$ with some norm $|.|$. We denote by C^0 the space of bounded continuous functions from $\mathcal{R}$ to $\mathcal{R}^n$ with norm $||x||_{C^0} = \sup_{t\in\mathcal{R}} |x(t)|$. Similarly, C^1 is the space of continuously differentiable functions from $\mathcal{R}$ to $\mathcal{R}^n$ with finite norm $||x||_{C^1} = \max(||x||_{C^0}, ||\dot{x}||_{C^0})$.

Theorem 1: An equilibrium x_0 of $\dot{x} = f(x)$ is hyperbolic iff the linear operator $L : C^1 \to C^0$, $L\xi = \dot{\xi} - Df_{x_0}\xi$, is invertible.

Theorem 2: A periodic solution $x_0(t)$ of a period T system $\dot{x} = f(x,t)$, is hyperbolic iff the linear operator $L : C^1 \to C^0$, $L\xi = \dot{\xi} - Df_{x_0(t),t}\xi$, is invertible.

Proofs can be found in [SS, Cop].

Now we are ready to generalise the definition of hyperbolicity to solutions of dynamical systems $\dot{x} = f(x,t)$ with other time-dependence. We suppose that f is differentiable in x and both f and the x-derivative Df are continuous in (x,t).

Definition: A solution $x_0(t)$ of a dynamical system $\dot{x} = f(x,t)$ is *hyperbolic* if the linear operator $L : C^1 \to C^0$, $L\xi = \dot{\xi} - Df_{x_0(t),t}\xi$, is invertible.

This formulation of hyperbolicity goes back to Massera and Schäffer [MS], who called it "exponential dichotomy". In the case of time-periodic dynamical systems (usually formulated in discrete time, by reduction to iteration of the map from time 0 to time T), it is equivalent to the standard dynamical systems notion of *uniform hyperbolicity* for arbitrary bounded orbits, though the notion is usually applied to invariant sets containing more than one orbit. The following result can be found in [SS, Cop] again.

Theorem 3: A bounded orbit $x_0(t)$ of $\dot{x} = f(x,t)$ is hyperbolic iff there exist $C, K, \alpha > 0$ and, for each $t \in \mathcal{R}$, a splitting $TM_{x_0(t),t} = E^+_{x_0(t),t} \oplus E^-_{x_0(t),t}$ such that

$$v \in E^+_{x_0(t),t} \Rightarrow |D\phi_{t,s}v| \le Ce^{-\alpha(s-t)}|v| \text{ for } s \ge t,$$

$$v \in E^-_{x_0(t),t} \Rightarrow |D\phi_{t,s}v| \le Ce^{-\alpha(t-s)}|v| \text{ for } s \le t,$$

$$v^+ \in E^+_{x_0(t),t}, v^- \in E^-_{x_0(t),t} \Rightarrow |v^+ + v^-| \geq K \max(|v^+|, |v^-|),$$

where $\phi_{t,s}$ is the evolution from time t to time s.

Since C, K and α are chosen independently of t, the term "uniformly hyperbolic" is often used, but we do not consider any other type of hyperbolicity here. It is automatic that $E^\pm_{x_0(t),t}$ depend continuously on t (a standard exercise), in particular they have constant dimension. The case where $\dim E^- = 0$ is called *linearly attracting.*

Note that by our definition, an autonomous dynamical system (in continuous time) can not have any hyperbolic solutions other than equilibria, because if f is independent of t then the derivative of a solution is in the kernel of L. The standard approach to the autonomous case is to modify the definition of the splitting to $TM_{x_0(t),t} = E^+_{x_0(t),t} \oplus E^-_{x_0(t),t} \oplus \mathcal{R}f$, and the conclusions of the theory have to allow a time-reparametrisation. We will not treat that case in this paper, but see [CFS].

The key result of the theory of uniform hyperbolicity for our purposes is (e.g. [Cop]; see [Fr] for the discrete-time analogue):

Theorem 4: If x_0 is a hyperbolic solution for $\dot{x} = f_0(x,t)$, then for $f_\varepsilon, Df_\varepsilon$ close to f_0, Df_0, there is a locally unique solution x_ε of $\dot{x} = f_\varepsilon(x,t)$ near x_0. Furthermore, x_ε is hyperbolic and depends smoothly on ε (as smoothly as f_ε depends on its arguments), and $E^\pm_{x_\varepsilon(t),t}$ depend continuously on ε (in particular, they keep the same dimensions).

This can be proved by using the implicit function theorem (IFT) on the map $G : \mathcal{R} \times C^1 \to C^0$ defined by

$$G_\varepsilon[x](t) = \dot{x}(t) - f_\varepsilon(x(t), t). \tag{1}$$

Although most proofs of the IFT (e.g. [D]) provide estimates on the continuation, one can often do better by using the chain rule to obtain

$$\frac{dx_\varepsilon}{d\varepsilon} = L_\varepsilon^{-1} \frac{\partial f_\varepsilon}{\partial \varepsilon}(x_\varepsilon), \tag{2}$$

where $L_\varepsilon = DG_{\varepsilon,x_\varepsilon}$, and bounding its right hand side. Here and in the following, for a function $G : \mathcal{R} \times X \to Y$ with X, Y function spaces, we use DG to denote the derivative in X and $\frac{\partial G}{\partial \varepsilon}$ for the derivative in $\mathcal{R}$. This equation can be used to good effect to estimate

1. a range of ε for which the continuation x_ε is guaranteed to exist (we call this *persistence*), and

2. a range of ε for which x_ε exists and remains within some desired open neighbourhood of x_0, which we call a "safety region".

Furthermore, if x_0 is attracting then so is x_ε (in both the forward and pullback senses; for definitions see [CKS]), by extension of Theorem 4 to give local forward and backward contracting submanifolds tangent to $E^\pm$ (or usual stability proofs for linearly stable solutions).

Here is the basic strategy, though it can be refined if given more information about the system. If L_0 is invertible, let

$$K_0 = ||L_0^{-1}||^{-1} > 0, \tag{3}$$

where for a bounded linear map A from one Banach space X to another Y, $||A||$ is the "operator norm" $\sup\{||Ax||_Y : ||x||_X = 1\}$. Suppose that for (ε, η) in a neighbourhood N of $(0,0)$ there is $\beta(\varepsilon, \eta) < K_0$ such that for all $t \in \mathcal{R}$ and $x(t)$ within η of $x_0(t)$ we have

$$|Df_{\varepsilon,x(t),t} - Df_{0,x_0(t),t}| \leq \beta(\varepsilon, \eta). \tag{4}$$

Then standard perturbation theory of linear operators, e.g. [K] implies that L_ε is invertible as long as (ε, η) remains in N, where η now denotes $||x_\varepsilon - x_0||_{C^0}$. Furthermore, the standard inequality $||B^{-1}||^{-1} \geq ||A^{-1}||^{-1} - ||B-A||$ (provided the right hand side is positive) implies

$$||L_\varepsilon^{-1}||^{-1} \geq K_0 - \beta(\varepsilon, \eta). \tag{5}$$

Suppose also that $\alpha(\varepsilon, \eta)$ is such that

$$|\frac{\partial f_\varepsilon}{\partial \varepsilon}(x,t)| \leq \alpha(\varepsilon, \eta), \tag{6}$$

for all (x,t) such that $|x - x_0(t)| \leq \eta$. Then, noting that L_ε is a bounded linear map from C^1 to C^0, (2) implies that

$$||\frac{dx_\varepsilon}{d\varepsilon}||_{C^1} \leq ||L_\varepsilon^{-1}||^{-1} \; ||\frac{\partial f_\varepsilon}{\partial \varepsilon}(x_\varepsilon)||_{C^0}. \tag{7}$$

In particular, $||\frac{dx_\varepsilon}{d\varepsilon}||_{C^0}$ is bounded by the same quantity. So x_ε can be continued at least as long as the solution $\eta(\varepsilon)$ of

$$\frac{d\eta}{d\varepsilon} = \frac{\alpha(\varepsilon, \eta)}{K_0 - \beta(\varepsilon, \eta)} \tag{8}$$

starting from $(\varepsilon, \eta) = (0,0)$ remains in the region $\beta(\varepsilon, \eta) < K_0$. Taking the supremum $\varepsilon_{\max}$ of the allowed values of ε guarantees continuability for all ε up to (though not including) $\varepsilon_{\max}$, and hence a region of functions f_ε for persistence of x_0. Furthermore,

$$||x_\varepsilon - x_0||_{C^0} \leq \eta(\varepsilon), \tag{9}$$

and so if the safety region $||x - x_0||_{C^0} < \eta_{\max}$ is desired, it suffices to keep ε below the first value such that $\eta(\varepsilon) = \eta_{\max}$ or $\varepsilon = \varepsilon_{\max}$.

The astute reader will notice that the same bound (9) applies to the C^1 norm of the response, by simply defining $\eta = ||x_\varepsilon - x_0||_{C^1}$ and using $||x||_{C^0} \leq ||x||_{C^1}$.

The term K_0 defined in eq. (3) is easily estimated when x_0 is an equilibrium of an autonomous vector field f_0. One simply finds the Green function g (the impulse response of the linearised system, i.e. the unique bounded solution of the distributional differential equation $\dot{g} = Df_{x_0} g + \delta$) and then for $v \in C^0$,

$$(L_0^{-1} v)(t) = \int_{-\infty}^{\infty} g(t-s) v(s) ds, \tag{10}$$

from which

$$||L_0^{-1}|| = \max \left(\int_{-\infty}^{\infty} |g(s)| ds, 1 + (\int_{-\infty}^{0} + \int_{0}^{\infty}) |\dot{g}(s)| ds \right). \tag{11}$$

Similarly, suitable functions α and β are often easily obtained. In the simplest case, where the perturbation is just an additive force u:

$$\dot{x} = f_0(x,t) + u(t) \tag{12}$$

with f_0 Lipschitz in x with Lipschitz constant C, then we can take $\beta = C\eta$ and writing

$$\varepsilon = ||u||_{C^0}, \tag{13}$$

we have $||\frac{\partial f_\varepsilon}{\partial \varepsilon}||_{C^0} = 1$. So the differential equation (8) becomes

$$\frac{d\eta}{d\varepsilon} = \frac{1}{K_0 - C\eta}. \tag{14}$$

Thus by integration we have a unique continuation $x[u]$ which satisfies

$$||x[u] - x_0||_{C^1} \leq \eta(\varepsilon) = \left(K_0 - \sqrt{K_0^2 - 2\varepsilon C} \right) / C,$$

the smaller solution of

$$K_0\eta - \frac{1}{2}C\eta^2 = \varepsilon, \tag{15}$$

as long as

$$\varepsilon < \frac{K_0^2}{2C}. \tag{16}$$

The idea is easily generalised to functional differential equations $\dot{x} = F[x]$ for C^1 maps $F : C^1 \to C^0$, for example differential delay equations like $\dot{x}(t) = f(x(t), x(t-1))$ [HV]. One simply defines $G : C^1 \to C^0$ by $G[x] = \dot{x} - F[x]$ and replaces (4) and (6) by $||DF_{\varepsilon,x} - DF_{0,x_0}||_{C^0} \leq \beta(\varepsilon, \eta)$ and $||\frac{\partial F}{\partial \varepsilon}||_{C^0} \leq \alpha(\varepsilon, \eta)$ for $||x - x_0||_{C^1} \leq \eta$. We will not pursue this generalisation here, though we will use this more general setting in Section 8.

3 Improved estimates from Equivalent norms

In practice the above strategy may give greatly suboptimal results. For example, the desired safety criterion might involve only $||x||_{C^0}$ and not $||\dot{x}||_{C^0}$, in which case larger safe domains would be found by putting less weight on the derivative in the C^1 norm. The standard C^1 norm is in any case not dimensionally consistent, so the estimates depend on the unit of time, and one should at least use a *weighted* C^1 norm like $\max(||x||_{C^0}, T||\dot{x}||_{C^0})$ for some $T > 0$. Normally T should be chosen to be a typical timescale for the system, or to reflect the weight one wants to give to $\dot{x}$ in a safety criterion, but for a C^0 safety criterion, the optimum results over T are achieved by taking any T small enough that estimates of $||L^{-1}||$ are dominated by the size of the Green function rather than its derivative. In some circumstances, other equivalent norms may be more appropriate, like $(||x||_{C^0} + T||\dot{x}||_{C^0})$. The formulae of the preceding section remain correct for any choice of norm on C^1 with the property $||x||_{C^0} \leq ||x||_{C^1}$. Similarly, the concept of persistence depends only on the topology defined by the C^1 norm, not the norm itself, so better persistence results might be obtained by replacing it by a well chosen equivalent norm.

A second loss of accuracy results from the bound (5), which although near optimal for small ε, η, may be too pessimistic for larger values. Better estimates can be obtained by some tricks, as we shall illustrate in Section 5.

As a third and more profound example, a weakly damped oscillator is sensitive to forcing near its natural frequency (or frequencies) but much less sensitive to forcing at distant frequencies. Thus in the case of additive forcing, for example, the true region of safe forcing functions might be far from being shaped like a C^0-ball. Mathematically, the point is that $L_\varepsilon^{-1}\frac{\partial f_\varepsilon}{\partial \varepsilon}$ might be much smaller than estimated in (7).

To take this third point into account, we propose the use of an *adapted* norm on C^0 in place of the standard C^0-norm. Specifically, we adapt the norm to the linearisation L_0 (assumed to be invertible), by defining

$$||u||_* = ||L_0^{-1}u||_{C^1}. \tag{17}$$

For example, if x_0 is a linearly attracting equilibrium of an autonomous system, with (matrix) impulse response function g, then $L_0^{-1}u = g * u$, the convolution of u with g:

$$(g * u)(t) = \int_{-\infty}^{t} g(t-s)u(s)ds, \tag{18}$$

and then $||u||_* = ||g * u||_{C^1}$. It is easy to show that if $L_0 : C^1 \to C^0$ is invertible then $||.||_*$ is equivalent to $||.||_{C^0}$.

The advantage of the $*$-norm is that $L_0 : C^1 \to C^0_*$ (C^0_* meaning the space C^0 with the $*$-norm) is an isometry: $||L_0\xi||_* = ||\xi||_{C^1}$, so in the early stages of the continuation only a small amount is lost by the estimate

$$||L_\varepsilon^{-1}\frac{\partial f_\varepsilon}{\partial \varepsilon}||_{C^1} \leq ||L_\varepsilon^{-1}||_{C^0_* \to C^1}||\frac{\partial f_\varepsilon}{\partial \varepsilon}||_*. \tag{19}$$

The effect on L of the change in Df then needs estimating as an operator from C^1 to C^0_*:

$$||L_\varepsilon - L_0||_{C^1 \to C^0_*} = \sup_{\xi \neq 0} \frac{||L_0^{-1}((Df_\varepsilon - Df_0)\xi)||_{C^1}}{||\xi||_{C^1}} \leq \beta(\varepsilon, \eta) \tag{20}$$

This may be difficult to do accurately, but β can always be taken to go to zero as $\varepsilon, \eta \to 0$. The end result is

$$||\frac{dx}{d\varepsilon}||_{C^1} \leq (1 - \beta(\varepsilon, \eta))^{-1}||\frac{\partial f}{\partial \varepsilon}||_*, \tag{21}$$

and this bound is asymptotically optimal as $\varepsilon, \eta \to 0$. Furthermore, $\frac{dx}{d\varepsilon}$ is close to constant for small ε, so the result of integrating this inequality is also asymptotically optimal.

Even for ε, η large, the adapted norm is often beneficial. For additive forcing, for example, we can redefine $\varepsilon = ||u||_*$ and then we have a continuation $x[u]$ with $||x[u] - x_0||_{C^1} \leq \eta(\varepsilon)$, the smaller solution of $\eta - \frac{1}{2}C\eta^2 = \varepsilon$, as long as $\varepsilon < \frac{1}{2C}$. This gives much larger regions of u for continuability of x_0 or for given safety limit $\eta_{\max}$, as will be demonstrated for some example systems in the following sections.

The idea of adapted norm can be refined if desired. Firstly, one can combine it with the effect of changing C^1 norms as in the first part of this section. Secondly, there is no need to use exactly the operator L_0^{-1} to define the $*$-norm, and it could be allowed to vary with x in order to adapt to the form of the nonlinearity of the system, though this complicates the calculations.

One might ask how close to optimal one could get by further refinements of IFT estimates. The IFT gives unique continuability as long as DG remains invertible; the continuation satisfies (2). Thus for problem 1 one can push ε arbitrarily close to $\varepsilon_c = \sup\{\varepsilon : x_{\varepsilon'}$ is defined and $DG_{\varepsilon',x_{\varepsilon'}}$ is invertible for all $\varepsilon' \in [0,\varepsilon)\}$, by sufficiently accurate approximation of the right hand side of (2) and its solution.

For problem 2, the solution x_ε might exit the desired safety region S before ε_c, but sufficiently accurate approximation of (2) will again show x_ε to remain inside S for all $\varepsilon <$

$$\varepsilon_s = \sup\{\varepsilon : x_{\varepsilon'} \text{ is defined and in } S \text{ and } DG_{\varepsilon',x_{\varepsilon'}} \text{ is invertible } \forall \varepsilon' \in [0,\varepsilon)\}.$$

The main defect of our approach is that invertibility of $DG_{\varepsilon,x_\varepsilon}$ might be lost for some ε but with x_ε remaining in S up to that ε and there might be a nearby attractor in S for larger ε (e.g. a "safe" bifurcation in the language of [TSU]), so safety would still be preserved even though the IFT could not prove it. Nevertheless, we believe our approach will be a useful tool, and we now give some examples of its application.

4 Additively forced damped oscillators

We treat now a specific class of examples, namely additive forcing of a stable equilibrium of a damped second-order oscillator

$$\ddot{x} + \gamma\dot{x} + V'(x) = u(t), \tag{22}$$

with $\gamma > 0$, $V'(0) = 0$, $V''(0) = \omega^2 > 0$. This is a commonly studied class of systems, e.g. [STUL]. For example, as a simple model of ship capsize one could take

$$V(x) = A\cos x - B\cos 2x, \tag{23}$$

with $0 < A < 4B$. An even simpler commonly used model is

$$V(x) = ax^2 - bx^4, \tag{24}$$

which approximates (23) near $x = 0$ ($a = 2B - A/2$, $b = 2B/3 - A/24$), and for which $\omega = \sqrt{2a}$.

We could transform (22) into a coupled system of two first-order differential equations to fit it into the framework of the previous sections, but rather than doing this, it is convenient to extend the approach of this paper to the map $G_u : C^2 \to C^0$ defined by

$$G_u[x] = \ddot{x} + \gamma\dot{x} + V'(x) - u. \tag{25}$$

For a given $u \in C^0$, a function $x \in C^2$ is a solution of (22) iff $G_u[x] = 0$. It is hyperbolic iff $DG_{u,x} : C^2 \to C^0$, given by

$$DG_{u,x}\xi = \ddot{\xi} + \gamma\dot{\xi} + V''(x)\xi, \tag{26}$$

is invertible. At the linearly attracting equilibrium $x = 0$ of the unforced case ($u = 0$), let us write $L_0 = DG_{0,0}$. The impulse response (bounded solution of $L_0 g = \delta$) is

$$g(t) = 0 \text{ for } t < 0, \frac{e^{\lambda_2 t} - e^{\lambda_1 t}}{\lambda_2 - \lambda_1} \text{ for } t \geq 0, \tag{27}$$

where λ_j are the roots of $\lambda^2 + \gamma\lambda + \omega^2 = 0$, assumed for simplicity to be distinct (the case of a double root is easily treated, and the results agree with the limit as $\lambda_2 - \lambda_1 \to 0$, e.g. $g(t) = te^{\lambda t}$ for $t \geq 0$).

For each $v \in C^0$, the unique solution of $L_0\xi = v$ in C^2 is $\xi = g * v$. Its derivatives are easily determined to be

$$\begin{aligned} \dot{\xi} &= \dot{g} * v \\ \ddot{\xi} &= v - \gamma\dot{g} * v - \omega^2 g * v, \end{aligned} \tag{28}$$

(note that $g(0) = 0$). Let us use the weighted C^2 norm

$$||x||_{C^2} = \max(||x||_{C^0}, T||\dot{x}||_{C^0}, T^2||\ddot{x}||_{C^0}), \tag{29}$$

where $T > 0$ is either specified (representing the relative weights to be given to tilt of the ship, angular velocity and acceleration) or is to be determined

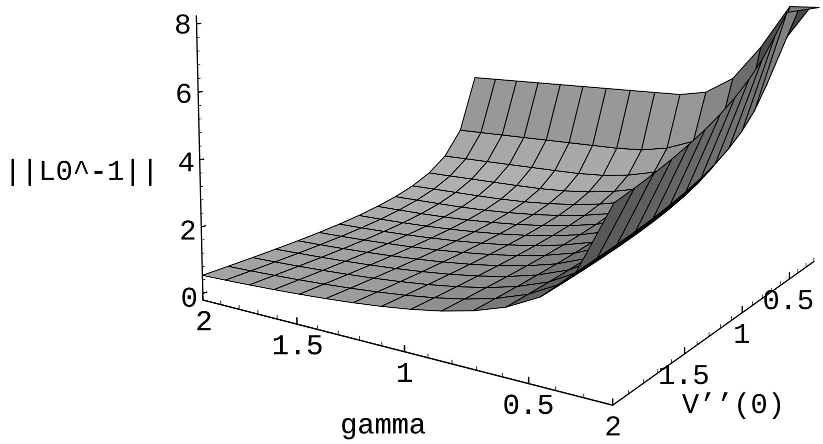

Figure 1: Bound (30) on $||L_0^{-1}||$ as a function of damping γ and force constant $V''(0) = \omega^2$.

to maximise the set of allowed forcing function for a safety criterion of the form $||x||_{C^0} \leq \eta$. Thus we have the estimate

$$\begin{aligned} ||L_0^{-1}|| &= \sup_{||v||=1} ||g * v||_{C^2} \\ &\leq \max\left(||g||_1, T||\dot{g}||_1, T^2(1 + \gamma||\dot{g}||_1 + \omega^2||g||_1)\right), \end{aligned} \tag{30}$$

where for a function ψ, $||\psi||_1 = \int_{-\infty}^{\infty} |\psi(t)|dt$. Figure 1 shows how this bound on $||L_0^{-1}||$ varies with damping γ and force constant ω^2 for $T = 1$ ($T \neq 1$ can be read off by reinterpreting the axes as $T\gamma, T^2\omega^2$ and $T^{-2}||L_0^{-1}||$).

The implicit function theorem now guarantees a locally unique continuation $x_u \in C^2$ for all $||u||_{C^0}$ small enough.

We compute a neighbourhood of $0 \in C^0$ of the form $||u||_{C^0} \leq \varepsilon(\eta)$ for which there is a unique bounded response in the safety region $||x||_{C^2} \leq \eta$. Denote $L_u = DG_{u,x_u}$. Then

$$L_u - L_0 = V'' \circ x_u - \omega^2, \tag{31}$$

so let β be any function such that

$$|V''(x) - \omega^2| \leq \beta(s) \tag{32}$$

for all $|x - x_0| \leq s$ (without loss of generality, $\beta(0) = 0$ and β is continuous). Then denoting $K_0 = ||L_0^{-1}||^{-1}$ as before, we obtain

$$||L_u^{-1}|| \leq (K_0 - \beta(\eta))^{-1}, \tag{33}$$

and hence for $u = \varepsilon v$ with $||v||_{C^0} = 1$,

$$||\frac{dx}{d\varepsilon}||_{C^2} \leq ||L_u^{-1}|| \leq (K_0 - \beta(\eta))^{-1} \tag{34}$$

as long as $\beta(\eta) < K_0$. We deduce that ε can be taken to be

$$\varepsilon(\eta) = \int_0^\eta (K_0 - \beta(s))\ ds. \tag{35}$$

For the particular potential (24), we can take

$$\beta(s) = 12bs^2, \tag{36}$$

and hence we can take

$$\varepsilon(\eta) = K_0\eta - 4b\eta^3 \tag{37}$$

as long as

$$\eta < \sqrt{K_0/12b}. \tag{38}$$

The results are shown in Figure 2 for $a = 1/2, b = 1/4$ and $\gamma = 0.5$.

If as safety criterion we choose only $||x||_{C^0} \leq \eta$, then we are free to optimise the result (37) over $T > 0$ (K_0 depends on T). The optimum is achieved for any T small enough that the first term of (30) dominates. The results of this optimisation are shown as a function of γ for $a = 1/2$, $b = 1/4$ (all other positive choices are related by scaling) in Figure 3. In particular, by taking the limiting value $\sqrt{K_0/12b}$ of η and optimising over T, we obtain persistence (in this context meaning locally unique continuability in C^2) of the bounded solution for all forcing functions u with

$$||u||_{C^0} < \frac{\sup_{T>0} K_0^{3/2}}{3\sqrt{3b}}, \tag{39}$$

shown in Figure 4 for $a = 1/2, b = 1/4$.

Note that since the dimension of the unstable subspace for the solution can not change during continuation, the resulting solution is still linearly attracting. Also, although persistence does not imply that the solution remains

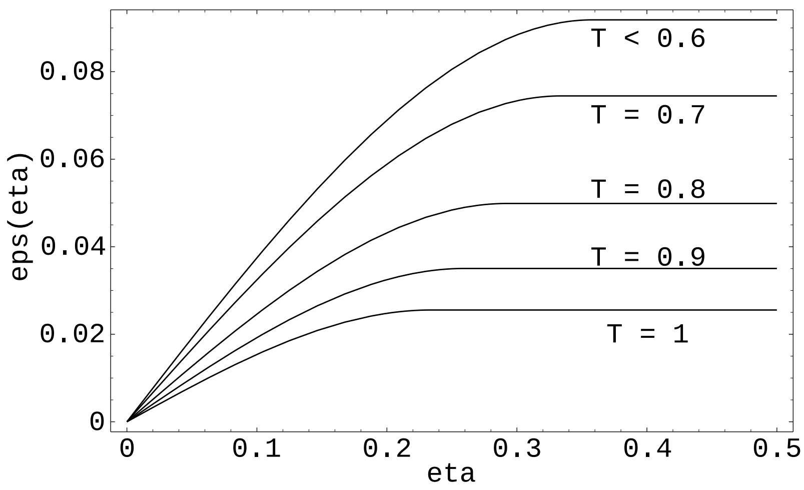

Figure 2: Plot of ε against η for different values of T. Note that the result is independent of T if small enough that the first term of eq. (30) dominates.

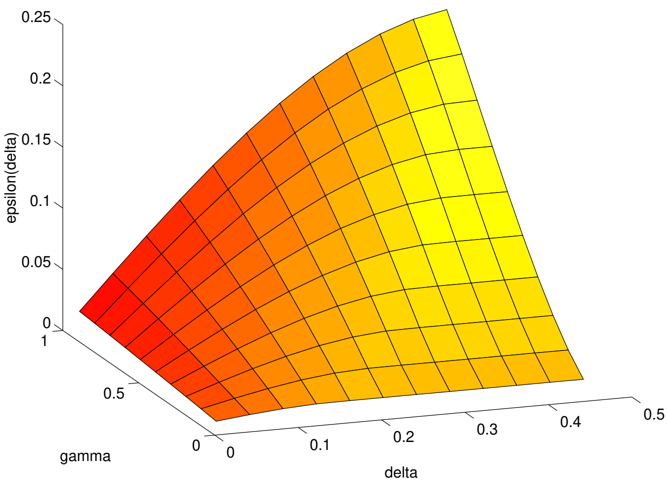

Figure 3: Plot of ε as a function of γ and η.

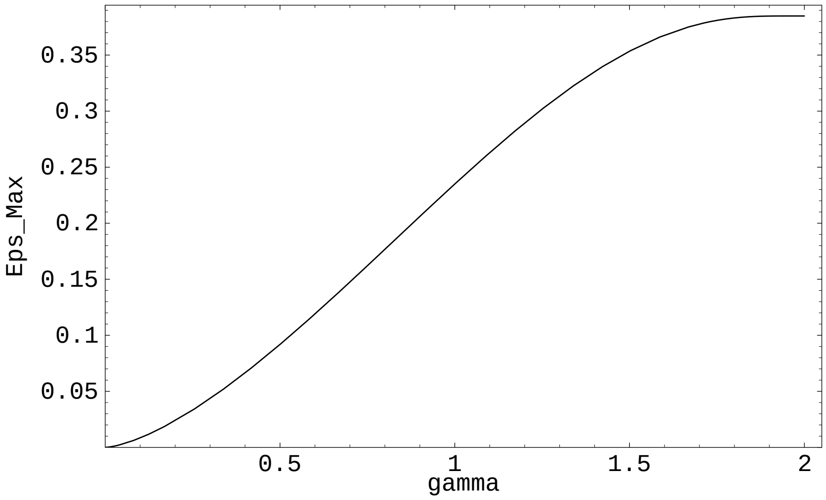

Figure 4: Plot of $\varepsilon_{\max}$ as a function of γ.

in the potential well around 0, the above result is produced using the bound (38), which can be shown to keep the solution between the inflection points within the well.

The result of Figure 4 can be slightly improved by using a more sophisticated estimate than (33), to be introduced in Section 5. For small damping, however, a much more significant improvement comes from using an adapted norm, as explained in Section 3, which permits much larger forcing functions provided they are sufficiently off resonance.

The adapted norm we use on forcing functions is given by

$$||u||_* = ||L_0^{-1}u||_{C^2} = ||g * u||_{C^2}. \tag{40}$$

which measures C^0 functions by the response they induce in the linearisation of the unperturbed system.

Taking β as in (32) and $u = \varepsilon v$ with now $||v||_* = 1$, we can use the estimate

$$||\frac{dx}{d\varepsilon}||_{C^2} \leq ||L_u^{-1}||_{C^0_* \to C^2} \leq (1 - K_0^{-1}\beta(\eta))^{-1}. \tag{41}$$

Hence we now obtain the bound

$$\varepsilon(\eta) = \int_0^\eta (1 - K_0^{-1}\beta(s))ds, \tag{42}$$

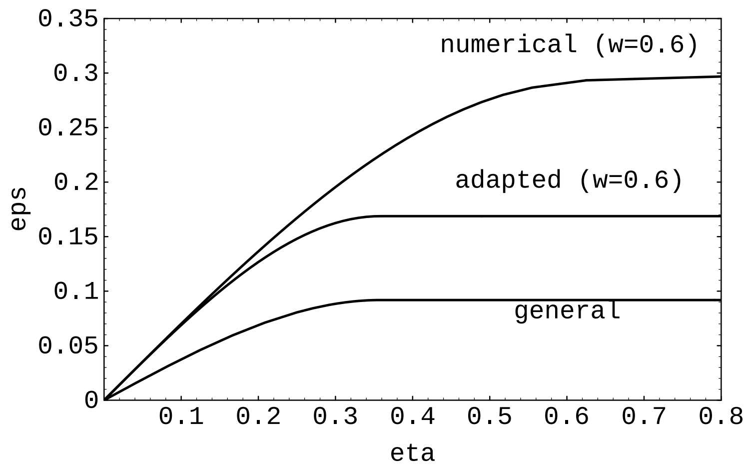

Figure 5: Estimates of $\varepsilon(\eta)$ when $u(t) = \sin(\omega t)$, $\omega = 0.6$, $\gamma = 0.5$.

again as long as $\beta(\eta) < K_0$. The advantage here is that forcing functions away from resonance have very small size in the adapted norm.

For the potential function given by (24), we have $\beta(s) = 12bs^2$, so we can take

$$\varepsilon(\eta) = \eta - 4K_0^{-1}b\eta^3, \tag{43}$$

again as long as $\eta < \sqrt{K_0/12b}$.

The simplest class of forcing functions we can treat this way is periodic forcing of the form $u = \varepsilon \sin \omega t$. In this case, if we weight our C^2-norm towards the C^0-norm by taking T sufficiently small, we have

$$||u||_\infty = \sup_{t\in[0,2\pi/\omega)} |\varepsilon \int_0^\infty g(s)\sin(\omega(t-s))ds|, \tag{44}$$

which is easily computed.

Figure 5 shows three curves for $\varepsilon(\eta)$ using the original estimates, these adapted estimates and a numerical simulation, for $\omega = 0.6$ and $\gamma = 0.5$. Figure 6 shows $\varepsilon(\eta)$ for a range of frequencies.

For additive forcing with a given bound, it is well known that the control strategy which induces the largest response is of *bang-bang* type, that is, the control always takes the maximum amplitude but switches direction when the state crosses a "switching locus" (e.g. [Son]). The switching locus is

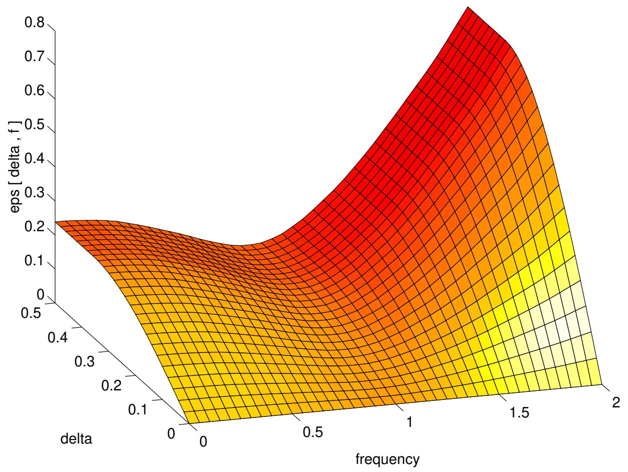

Figure 6: Estimates of $\varepsilon(\eta)$ when $u(t) = \sin \omega t$, as a function of η (denoted 'delta' in the figure) and ω (denoted 'frequency'), using the adapted norm procedure, with $\gamma = 0.5$.

easily obtained in our case. The slope of the trajectory in the phase-plane is $d\dot{x}/dx = (-\gamma\dot{x} - V'(x) + u(t))/\dot{x}$. In the upper half-plane, $\dot{x} > 0$, so increasing $d\dot{x}/dx$ gives bigger solutions. Conversely, for $\dot{x} < 0$, decreasing $d\dot{x}/dx$ gives bigger solutions. So the switching locus is $\dot{x} = 0$ and the bang-bang control given by

$$u^*(t) \;=\; \begin{cases} \varepsilon\ , & \dot{x} > 0 \\ -\varepsilon\ , & \dot{x} < 0 \end{cases} \tag{45}$$

is the worst case forcing function.

In order to check how accurate our safety criteria are we numerically integrated our nonlinear oscillator with a variety of forcing functions including the optimal bang-bang control given above and made a comparison.

Figure 7 shows the response of the system for various forcing functions along with the hyperbolicity estimates of the response given by (37). The periodic forcing frequency for both the sinusoidal and 'square wave' forcing is chosen so as to give the biggest response ($\omega \approx 0.6$) but it is clearly less efficient at producing large motions than the bang-bang control, particularly when the forcing amplitude ε exceeds about 0.23, the first value at which bang-bang control leads to escape from the well.

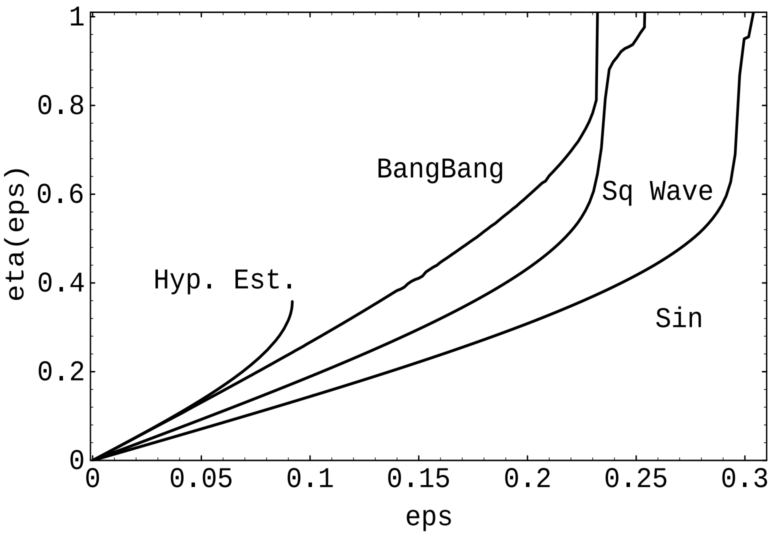

Figure 7: Comparison of safety estimates obtained from persistence theorem and from numerical simulation.

As already discussed, the hyperbolicity estimates using the adapted norm are asymptotically optimal for small forcing amplitudes (provided that also the derivative is weighted sufficiently weakly in the norm on C^1). This is borne out in the figures where it is clear that the slopes of the hyperbolicity estimates and the numerics for the bang-bang control are equal at the origin.

5 Parametrically and additively forced non-linear oscillators

We consider now the damped oscillator with parametric as well as additive (external) forcing:

$$\ddot{x} + \gamma\dot{x} + (1 + v(t))V'(x) = u(t), \tag{46}$$

with $u, v \in C^0$ and $\gamma > 0$, $V'(0) = 0$, $V''(0) = \omega^2 > 0$ as before.

In [TRS], some detailed physical arguments are given to derive a ship capsize equation of this form. Here we consider the realistic assumption that the perturbations are bounded and continuous but otherwise arbitrary. The operator we consider is

$$G_{u,v}[x] = \ddot{x} + \gamma\dot{x} + (1 + v(t))V'(x) - u. \tag{47}$$

In this case

$$DG_{u,v,x}\xi = \ddot{\xi} + \gamma\dot{\xi} + (1 + v(t))V''(x)\xi, \tag{48}$$

Since $L_0 = DG_{0,0,0}$ is the same constant coefficient operator as in the previous section, it is invertible. Equivalently, the linearly attracting solution, $x = 0$ for the unforced ($u = v = 0$) system is hyperbolic. So using the implicit function theorem, for small enough $u, v \in C^0$, we have a locally unique continuation $x_{u,v} \in C^2$.

Let $L_{u,v}$ denote $DG_{u,v,x_{u,v}}$. Then

$$L_{u,v} - L_0 = V''(x_{u,v})(1 + v(t)) - \omega^2 \tag{49}$$

Let $K_0 = ||L_0^{-1}||^{-1}$ as before. Define $\varepsilon_1 = ||u||_{C^0}$, $\varepsilon_2 = ||v||_{C^0}$ and let $\varepsilon = \binom{\varepsilon_1}{\varepsilon_2}$. Suppose also that β is any function such that $||L_{u,v} - L_0|| \leq \beta(\eta, \varepsilon_2)$ whenever $||x||_{C^2} \leq \eta$. We then have the estimate

$$||L_{u,v}^{-1}|| \leq (K_0 - \beta(\eta, \varepsilon_2))^{-1} \tag{50}$$

which is valid while $\beta(\eta, \varepsilon_2) < K_0^{-1}$. While continuable, we have the differential equation

$$\frac{dx_{u,v}}{d\binom{u}{v}} = -L_{u,v}^{-1} \circ (V'(x_{u,v}), -I) \tag{51}$$

Taking $\alpha(\eta)$ such that $||V'(x_{u,v})|| \leq \alpha(\eta)$, we can use

$$\frac{d\eta}{d\varepsilon} \leq (K_0 - \beta(\eta, \varepsilon_2))^{-1}(\alpha(\eta), 1) \tag{52}$$

to provide safety estimates.

For the specific potential function $V(x) = ax^2 - bx^4$ we have

$$\alpha(\eta) = 2a\eta + 12b\eta^2 \tag{53}$$

and

$$\beta(\eta, \varepsilon_2) = (1 - \varepsilon_2)12b\eta^2 + 2a\varepsilon_2 \tag{54}$$

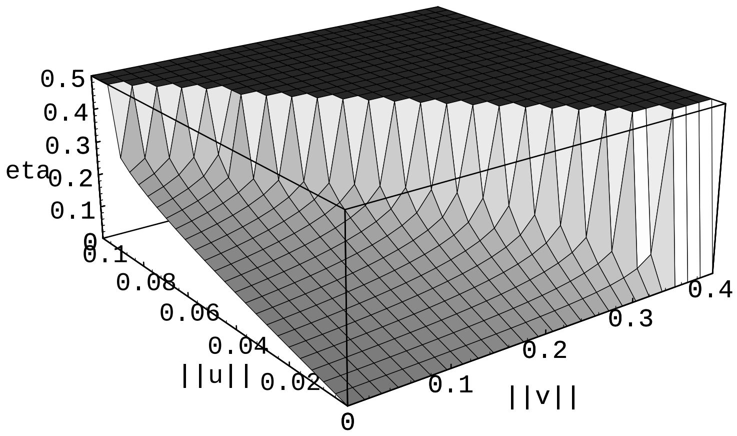

Figure 8: Safety criteria for the response of the parametrically and externally driven capsize equation. $\gamma = 1/2, a = 1/2, b = 1/4$

hence we can find $\eta(\varepsilon)$ by integrating (52) above. We can continue while

$$\eta^2 < (K_0^{-1} - 2a\varepsilon_2)((1-\varepsilon_2)12b)^{-1}. \tag{55}$$

Notice that small amplitude parametric forcing alone induces no response (though the response can become non-unique for larger amplitude). There has to be some external forcing as well. Thus $\eta\binom{0}{\varepsilon_2} = 0$ and we can obtain the surface $\eta(\varepsilon)$ by choosing lines of fixed ε_2 and integrating from $\varepsilon_1 = 0$. Figure 8 shows the surface for the capsize equation with $\gamma = 1/2, a = 1/2, b = 1/4$.

Another useful result related to parametric forcing is given in [RM] and we can use this approach to provide better upper bounds for $||L_u^{-1}||$ in the additively forced case. Although we have the general estimate $||L_u^{-1}|| \le (K_0 - \beta(\eta))^{-1}$, we can be more direct by treating L_u as a damped Hill operator with arbitrary bounded parametric excitation. In particular, we have

$$L_u\xi = \ddot{\xi} + \gamma\dot{\xi} + (V''(0) + p(t))\xi = 0 \tag{56}$$

where $p(t) = V''(x_u(t)) - V''(0)$.

When p is periodic this is a standard Hill's equation and, because of the viscous damping, for small enough excitation, $\xi = 0$ is an asymptotically stable equilibrium. However, larger excitations, especially near resonant frequencies, cause this solution to become unstable. In our case we need to

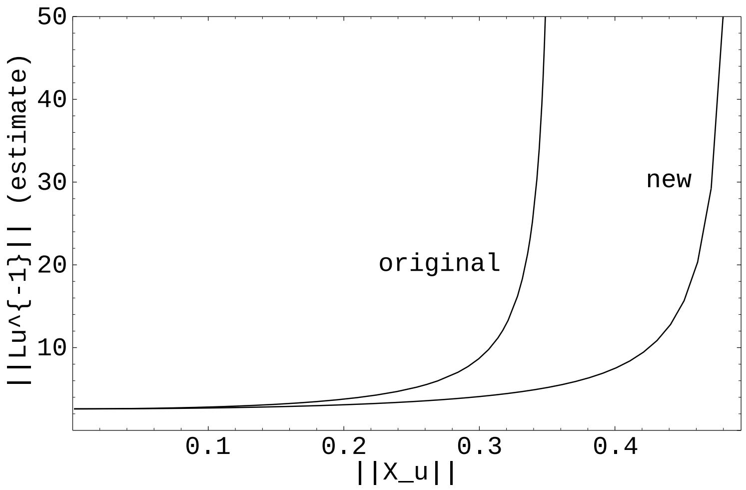

Figure 9: Upper bounds for $||L_u^{-1}||$ obtained from the original general estimate and an improvement based on the estimate of [RM] for the Green function. The parameters values used are $\gamma = 1/2, a = 1/2, b = 1/4$.

allow p to be aperiodic. For $||p||_\infty$ small enough, $\xi = 0$ remains asymptotically stable, $L_u\xi = 0$ has a bounded Green function and thus L_u^{-1} is bounded.

[RM] treat in depth the problem of estimating the Green function for p aperiodic and bounded. We do not give details here but the basic idea is to determine the worst case excitation and thus deduce a minimum decay rate $b^* < 0$ for solutions of $L\xi = 0$. From the bound $|\xi(t)| \leq Ne^{b^*(t-s)}|\xi(s)|$ we can easily bound the Green's function $g(t,s)$. For example, we can use the crude estimate $||g(t,\cdot)||_1 \leq -N/b^*$. although we can do better if careful. Figure 9 shows the improvement we gain from this method.

6 Phase-locked loops

As a second application, we consider phase-locked loops [Br]. The majority of these (those which use integrating filters) consist of oscillators which when unforced have an attracting periodic orbit (limit cycle) with a frequency ω_0, say. Under periodic forcing at a frequency ω near ω_0 (or more generally near a low order rational multiple of ω_0, but for simplicity we treat only the fundamental case here), the oscillator locks to the forcing, so the response is (asymptotically) periodic with frequency ω, and the phase of the response has a definite (asymptotic) relationship with the phase of the forcing. This is the familiar phenomenon of "frequency locking" or "phase locking" (no need to make a distinction here), epitomised in "Arnol'd tongues" in the dynamical systems literature, e.g. [KH].

In practical applications of phase-locked loops, however, the forcing is not exactly periodic, but it is desired that the phase of the response should track the phase of the forcing relatively closely. In particular, "cycle slipping" must be avoided.

The problem has several mathematically close relatives, a key one being that of whether an AC power generator can follow fluctuations in the electricity grid (e.g. [DK]).

To treat these problems, we start by considering an "ideal" forcing function, which is periodic with frequency near enough to ω_0 that the oscillator frequency locks to it. Then we ask for how large a perturbation from the ideal forcing function can we be sure that the response remains close.

Now if we required closeness of the response in C^1 then we would be limited to considering small amplitude modulation of the forcing, with fixed carrier frequency. But in reality we must also be prepared for frequency modulation: variations in the frequency of the forcing. Even if the frequency is fixed we must be prepared for large variations in the phase of the forcing.

To fit these types of perturbation into our scheme, we adopt the following trick. Suppose the function f_ε is such that there exists a C^1 function $t_\varepsilon : \mathcal{R} \to \mathcal{R}$ with derivative t'_ε C^0-close to 1 and $f_\varepsilon(x, t_\varepsilon(s))$ uniformly near to $f_0(x, s)$. This is the case for typical examples of frequency modulation or large phase shift discussed above. Then

$$\frac{dx}{ds} = f_\varepsilon(x, t_\varepsilon(s))t'_\varepsilon(s), \tag{57}$$

which is uniformly close to $f_0(x, s)$. So it is enough to consider C^0-small perturbations.

We suppose that for the ideal forcing f_0 there is an ("ideal") hyperbolic response x_0. Then the methods of the previous sections allow us to find a neighbourhood of forcing functions f_ε for which the response x_0 has a unique continuation x_ε, and in particular to find a subneighbourhood of forcing functions for which the response remains within a prescribed tube around the ideal response.

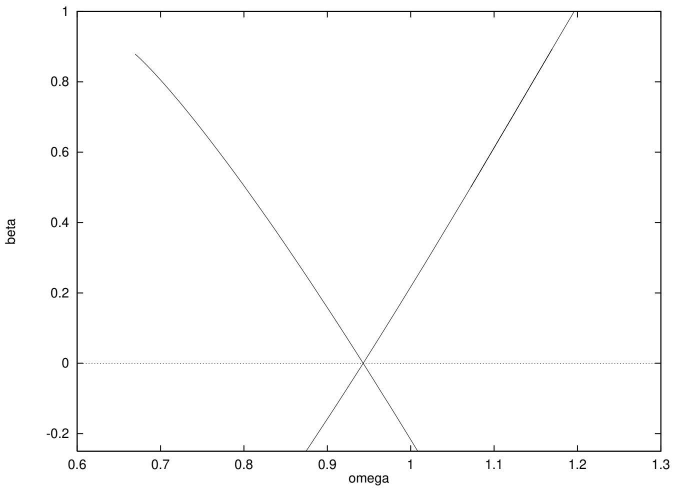

Figure 10: Region of (β, ω) space supporting an attracting $1:1$ phase-locked solution for the periodically forced van der Pol oscillator.

As an example, we consider a forced van der Pol oscillator

$$\begin{aligned} \dot{x} &= y - x^3/3 + x \\ \dot{y} &= -x + u(t) \end{aligned} \tag{58}$$

For the forcing function $u(t) = \beta \sin \omega t$, Figure 10 shows the region of (β, ω) space supporting an attracting $1:1$ phase locked solution.

We wish to determine a class of nearby forcing functions for which there remains a nearby linearly attracting response. Using the trick (57), we also allow forcing functions which become nearby after a time-reparametrisation. We write our forcing functions as

$$u(t) = C_\varepsilon(t) \sin \int_0^t \omega_\varepsilon(\tau) d\tau, \tag{59}$$

with ω_0 and C_0 constant. The time-reparametrisation we use is

$$\omega_0 ds = \omega_\varepsilon(t) dt, \tag{60}$$

so $u(t) = C_\varepsilon(t) \sin \omega_0 s$.

First we need to compute (numerically) the Green function $g(t, s) = (\xi, \eta)$ for the case $\varepsilon = 0$, i.e. the solution of the linearised equations about the periodic orbit $x_0(.)$ with delta-function force at time s:

$$\begin{aligned}\dot{\xi} &= \eta - x_0^2(t)\xi + \xi, \\ \dot{\nu} &= -\xi + \delta(t-s).\end{aligned} \tag{61}$$

By periodicity, it suffices to compute it for one cycle of s.

Next we have to choose a norm on the phase plane (x, y), for example $||(x, y)|| = \max(|x|, \kappa|y|)$, for some weight $\kappa > 0$ depending on the safety criterion, or which can be optimised over afterwards if only bounds on x are required. Then

$$||L_0^{-1}|| = \sup_{t\in[0,T]} \max\left(||\xi(t,.)||_1, \kappa||\nu(t,.)||_1, ||\dot{\xi}(t,.)||_1, 1 + ||\dot{\nu}(t,.)||_1\right), \tag{62}$$

where $T = 2\pi/\omega$. As usual we write K_0 for $||L_0^{-1}||^{-1}$.

For bounds α and β (cf. Section 2) we can use

$$\begin{aligned}\alpha(\varepsilon,\eta) &= \sup_{s\in\mathcal{R},|x-x_0(t_\varepsilon(s))|\le\eta} \\ &\quad \left(||\frac{\partial C_\varepsilon(t_\varepsilon(s))}{\partial\varepsilon}\sin\omega_0 s|| + ||f(x) + C_\varepsilon(t_\varepsilon(s))\sin\omega_0 s)||\ |\frac{\partial}{\partial\varepsilon}\frac{\omega_0}{\omega_\varepsilon(t_\varepsilon(s))}|\right), \\ \beta(\varepsilon,\eta) &= \sup_{t\in[0,T],|x-x_0(t)|\le\eta}\left(||Df_x - Df_{x_0(t)}|| + ||Df_x||\ |1 - \frac{\omega_0}{\omega_\varepsilon(t)}|\right).\end{aligned} \tag{63}$$

Then the standard differential equation (8) can be integrated to deduce a set of safe ε for given safety limit η.

We have treated a specific example in detail. Figure 11 shows a 1 : 1 phase-locked solution for the system

$$\begin{aligned}\dot{x} &= y - x^3/3 + x^2/2 + x/2 \\ \dot{y} &= -x + \kappa(y - u(t))\end{aligned} \tag{64}$$

with $u(t) = \cos(\omega t)$, $\omega = \pi/10$ and $\kappa = 0.4$ (we use the same κ in the C^1 norm). Although closely related to the standard van der Pol system, this one includes a symmetry breaking term to give it greater genericity. The chosen parameters give a system close to a period-doubling bifurcation, which occurs at approximately $\kappa = 0.419$. By applying the process described above we have calculated safety criteria and these are shown in figure 12. The largest change in forcing for which we can guarantee persistence is $\varepsilon = 0.0011$, for which we have $\eta(\varepsilon) = 0.021$.

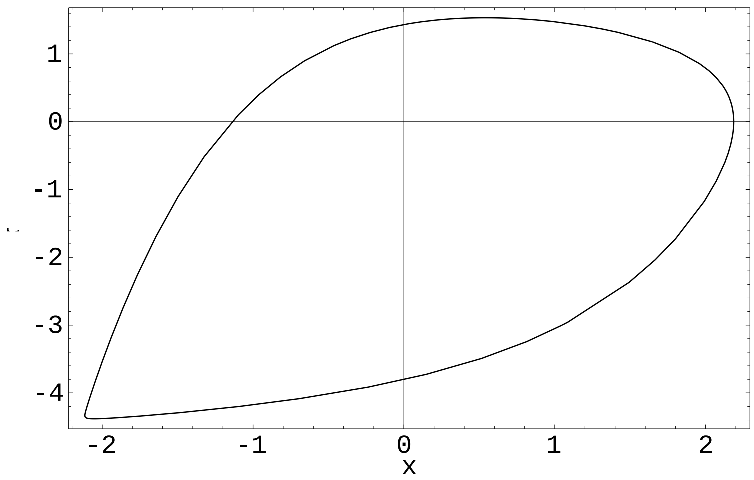

Figure 11: Periodic orbit close to period-doubling. $\kappa = 0.4$.

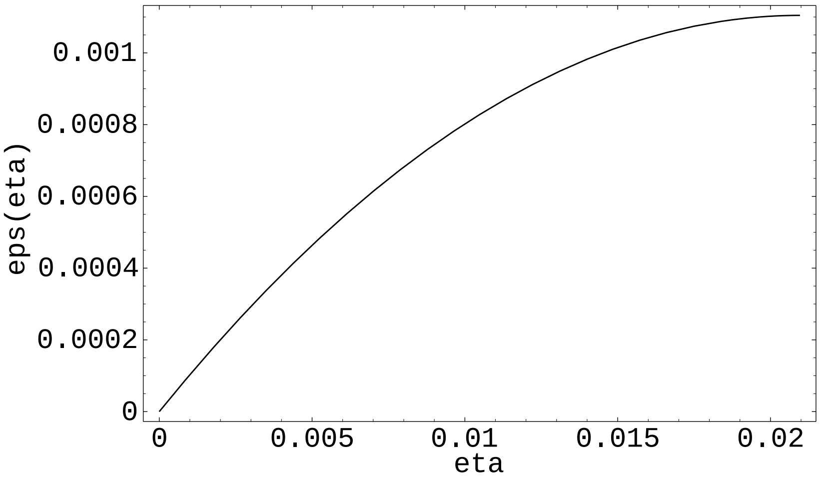

Figure 12: $\varepsilon(\eta)$ for a 1:1 phase-locked loop close to period-doubling bifurcation.

It would be interesting to test the method in a parameter regime near a phase-locking boundary, rather than a period-doubling bifurcation, e.g. $\kappa \approx 0.2, \omega \approx 0.9$, as cycle slipping is a more serious risk. Also, as in section 3, it would be good to weight the C^1-norm carefully and to adapt the C^0-norm to the linearisation about the ideal solution, which would make the method asymptotically optimal for small perturbation, as in the case of forcing an equilibrium.

7 Basins of attraction

A simple modification of the method provides inner estimates on the basin of attraction of a linearly attracting solution (whether an equilibrium, periodic or aperiodic), i.e. guaranteed neighbourhoods of attraction. One way to tackle this is to consider the special case of δ-function forcing, which can achieve any initial condition for a first-order system. This does not fit, however, into the framework as presented so far, because we have been requiring C^0 forcing, though in the next section we will extend to δ-function forcing. The problem of estimating a basin of attraction can be solved without this, however, simply by restricting the functional analysis to functions on a half-line (representing non-negative time only) and incorporating the initial condition into the function space.

For the ODE $\dot{x} = f(x,t)$, it is a standard result that if x_0 is a linearly attracting solution then it is also uniformly asymptotically stable. We say $(c, t_0) \in \mathcal{R}^n \times \mathcal{R}$ is in the *basin of attraction* of x_0 if the solution to $\dot{x} = f(x,t)$ with initial condition $x(t_0) = c$ converges to x_0 as $t \to +\infty$. Let $t_0 \in \mathcal{R}$. Let C^1_+ be the space of bounded C^1 functions defined on $[t_0, \infty)$ and let C^0_+ be the space of bounded C^0 functions defined on $[t_0, \infty)$. We can then use the implicit function theorem on the operator $G : \mathcal{R}^n \times C^1_+ \to C^0_+ \times \mathcal{R}^n$

$$G_c[x](t) = (\dot{x} - f(x,t) \ , \ x(t_0) - c) \tag{65}$$

to find a unique bounded continuation $x_c \in C^1_+, \ G_c[x_c] = 0,$ for $|x_0(t_0) - c|$ small enough. Since x_c is a root of G_c it solves the initial value problem $\dot{x} = f(x,t), \ x(t_0) = c$. We will show later that $x_c \to x_0$ as $t \to +\infty$.

The condition for continuation is that $L_0 = DG_{x_0(t_0),x_0} : C^1_+ \to C^0_+ \times \mathcal{R}^n$ given by $L_0\xi = \Big(\dot{\xi} - Df_{x_0}\xi \ , \ \xi(t_0)\Big)$ should be invertible. Since the spectrum of Df_{x_0} is in the open left-half plane this has a unique C^1_+ inverse given by the variation of parameters formula. For each $(\psi, b) \in C^0_+ \times \mathcal{R}^n$

$$L_0^{-1}(\psi, b) \;=\; X_0(t, t_0)\; b \;+\; \int_{t_0}^{t} X_0(t,\tau)\;\psi(\tau)\; d\tau, \tag{66}$$

where $X_0(t,\tau)$ is the evolution operator (or fundamental matrix solution) for $\dot{\xi} = Df_{x_0}\xi$.

Since x_0 is linearly attracting, $X_0(t,\tau)$ is exponentially decaying with constants independent of t_0. Let $K_H = \sup_{t\in[t_0,\infty)} |X_0(t,t_0)|$ and $K_P = \sup_{t\in[t_0,\infty)} \int_{t_0}^{t} |X_0(t,\tau)|\; d\tau$. Then the optimally weighted norm on $C^0_+ \times \mathcal{R}^n$ is $||(\psi,b)|| = ||\psi|| + K_H K_P^{-1}|b|$ and we obtain

$$||L_0^{-1}|| \leq K_P. \tag{67}$$

Thus we can obtain a good estimate of $||L_0^{-1}||$ independent of t_0. If we denote by $X_c(t,\tau)$, the evolution operator for $\dot{\xi} = Df_{x_c}\xi$, and write $L_c = DG_{c,x_c}$, we can use the differential equation

$$\frac{dx_c}{dc} \;=\; L_c^{-1} \circ (0, -I_{\mathcal{R}^n}) \;=\; X_c(\cdot, t_0) \tag{68}$$

to provide estimates of $\varepsilon_{\max}$ such that c is in the basin of attraction of x_0 for $|x_0(t_0) - c| < \varepsilon_{\max}$.

Let $\eta = ||x - x_0||$ as usual and assume we can find $\beta(\eta) \geq ||Df_x - Df_{x_0}||$. Then, in the absence of any other useful information about X_c, we can always use the bound $||L_c^{-1}|| \leq (K_P^{-1} - \beta(\eta))^{-1}$ for $\beta(\eta) < K_P^{-1}$. Now to find $\varepsilon_{\max}$ we just integrate

$$\frac{d\varepsilon}{d\eta} = K_H^{-1}(1 - K_P\beta(\eta)) \tag{69}$$

until we hit $\beta(\eta) = K_P^{-1}$.

Note that since all our estimates can be made independent of the choice of starting time t_0, we have *uniform* asymptotic stability and the basin of attraction we have found is uniform on the whole of $\mathcal{R}$.

It just remains to show that x_c while it exists converges to x_0 as $t \to +\infty$. This follows from the fact that L_c^{-1} is also exponentially decaying. In fact one can obtain estimates of decay rates for x_c if that is useful.

In many problems, large neighbourhoods of attraction can be found by constructing a Lyapunov function, but we feel that our approach could sometimes have advantages, in particular for aperiodic systems.

8 Discontinuous or distributional forcing

Our safety criteria, as developed so far, give regions of forcing functions for which the response remains within a given C^1-neighbourhood of the unperturbed response (C^2-neighbourhood in the case of section 4). Small C^1-norm is good for comfort, but for many applications, we need to know about only the C^0-norm of the response. Furthermore, insisting on C^1-limits on the response requires the forcing to be continuous, whereas small steps and even small impulses (δ-functions) and well behaved sequences of such, and almost all realisations of white noise, should also be alright.

So we propose to extend the problem of response to aperiodic forcing as follows. Let L^∞ be the space of equivalence classes of essentially bounded measurable functions $\{x : \mathcal{R} \to \mathcal{R}^n : \exists K \in \mathcal{R}\ s.t.\ |x(t)| \leq K\ a.e.\}$, under the equivalence relation $x \sim y$ if $x(t) = y(t)\ a.e.$, with the norm $||x||_\infty$ defined to be the infimum of such K. Let $W^{1,1}$ be the completion of the set of C^1 functions with respect to the norm

$$||\psi||_{1,1} = \int_{-\infty}^{\infty} (|\dot{\psi}(t)| + |\psi(t)|)dt. \tag{70}$$

For a Banach space B, the dual space B^* is the space of bounded linear functionals $u : B \to \mathcal{R}$ with norm $||u||_{B^*} = \sup_{b \in B \setminus 0} |u(b)|/||b||_B$.

Given a C^1 map $F : L^\infty \to (W^{1,1})^*$, e.g. $F[x](t) = f(x(t), t) + \varepsilon u(t)$ with u a piecewise continuous function plus some delta functions, define $G : L^\infty \to (W^{1,1})^*$ by $G[x] = \dot{x} - F[x]$, where the right hand side is interpreted in the weak sense:

$$\int G[x](t)\psi(t)dt = -\int x(t)\dot{\psi}(t) + F[x](t)\psi(t)\ dt$$

for all $\psi \in W^{1,1}$. So x is a (weak) solution of $\dot{x} = f(x,t)$ iff $G[x] = 0$. Thus zeroes of G qualify to be called solutions of the differential equation, even if not differentiable.

Then we wish to use continuation to deduce safety limits for the response of the solution x to changes in F. Firstly, it is easy to check that G is a C^1 map from L^∞ to $(W^{1,1})^*$ (this is especially clear if one uses the isomorphism $L^\infty = (L^1)^*$, so $x \mapsto \dot{x}$ is a linear map of norm 1). Secondly, at any C^1 solution x of a C^1 system which is hyperbolic in our sense, DG is also invertible as an operator from L^∞ to $(W^{1,1})^*$, because the inverse is still given by convolution with a Green function $g(.,s)$ which has bounded L^1-norm for each s.

Again, it would be better to adapt the norm on $(W^{1,1})^*$ (inherited from $W^{1,1}$) to the linearisation around the unperturbed solution, but we leave that refinement for future work.

The same idea can be applied to many other classes of forcing. One simply chooses an appropriate pair (A, B) of function spaces, differing by one degree of differentiability, such that $G : A \to B$, and checks that the IFT can be applied.

9 Concluding remarks

We have presented a general method to obtain safety criteria for aperiodically forced dynamical systems, i.e. sufficient conditions on forcing functions for the response of the system to remain within a given region. It is based on a non-autonomous version of the theory of uniform hyperbolicity. We have treated in detail the effects of forcing on an attracting equilibrium and of perturbation of periodic forcing on an attracting periodic orbit.

The continuation results work just as well for unstable hyperbolic solutions as for attracting ones, but practical application of the results would follow only if the unstable hyperbolic solutions belong to a set which is itself attracting, as for a uniformly hyperbolic strange attractor of an autonomous system. Using the ideas of the "shadowing lemma" (e.g. [KH]), analogous safety results could be obtained in this case. We remark that the shadowing lemma is an early result on aperiodic forcing of a uniformly hyperbolic dynamical system: the conclusion is that if the forcing is small and the solution remains close to a uniformly hyperbolic set then there is a uniformly close true orbit of the unforced system.

One criticism of our approach to safety criteria is that unique continuability of a trajectory with respect to forcing is too strong a requirement. For example, one might not mind if a periodically forced system goes through a supercritical Hopf bifurcation from an attracting periodic orbit to an attracting 2-torus, because even though the original orbit has more than one continuation (unless one insists on a periodic continuation with the same period), the response still remains nearby in phase space. Certainly the distinction between safe and catastrophic bifurcations (e.g. [TSU]) is an important one, but we feel that even though our approach gives pessimistic results in the case of a safe bifurcation, safe bifurcations are often a warning sign of disaster nearby and hence better avoided.

An alternative approach to derive safety criteria would be to use algorithms for time-optimal control to determine the minimum time to reach a given state from an undisturbed attracting equilibrium point by control functions satisfying a given bound ε (as described in an Appendix to [CK]). If the boundary of the safety region is unattainable by control functions bounded by ε (in which case the minimum time to reach it is declared infinite), then all such control functions are safe. But this approach requires infinite time calculations, so although it does not suffer from the above defect of our method, it does not give computable safety criteria.

As mentioned, our method can also be applied to many dynamical systems which are not of the standard vector field form, for example differential delay equations [HV]. It is anticipated that safety criteria can be designed similarly for many further types of situation. One example is when for no forcing the system has an (orbitally) linearly attracting periodic solution to which it is desired to remain close. A second example is when several different periodic forcing functions lead to nearly the same periodic response and it is desired to show that relatively sudden switches between the input functions preserve this property, e.g. quadrature phase-shift keying [UW]. The former would need the theory of "normal hyperbolicity" (e.g. [W]), and the latter would need an adaptation of the theory of shadowing.

Acknowledgements

We are grateful to Prof. Cullyer of the University of Warwick for stimulating the interest of RSM in this area in 1990, and to Allan McRobie for useful comments. The work was supported by an EPSRC Studentship for ZB and a British Council Alliance with the Université de Bourgogne, France, which made possible a visit of ZB while RSM was on leave there in 1994/95. We are grateful to the Laboratoire de Topologie for their hospitality that year.

References

[Arn] Arnold L, Random dynamical systems (Springer, 1998).

[BA] Bathrop NDP, Adams AJ, Dynamics of fixed marine structures (Butterworth-Heinemann, 3rd ed, 1991)

[Bi] Bishnani Z, Safety criteria for aperiodic dynamical systems, PhD thesis (Warwick, 2000)

[Br] Brennan PV, Phase locked loops – principles and practice (MacMillan, 1996).

[CKS] Cheban DN, Kloeden PE, Schmalfuss B, The relationship between pullback, forwards and global attractors of nonautonomous dynamical systems, J Dyn Sys Diff Eq 13 (2001) 185–213.

[CV] Chepyzhov VV, Vishik MI, Attractors of non-autonomous dynamical systems and their dimension, J Math Pures Appl 73 (1994) 279–333.

[CFS] Churchill RC, Franke J, Selgrade J, A geometric criterion for hyperbolicity of flows, Proc Am Math Soc 62 (1977) 137-143.

[CK1] Colonius F, Kliemann W, Some aspects of control systems as dynamical systems, J Dyn Diff Eq 5 (1993) 468–494.

[CK] Colonius F, Kliemann W, The dynamics of control (Birkhauser, 1999).

[Cop] Coppel WA, Dichotomies in stability theory, Lect notes in Math 629 (Springer, 1978).

[D] Dieudonné J, Foundations of modern analysis (Acad Press, 1969).

[DK] Dorning JJ, Kim JH, Bridging the gap between the science of chaos and its technological applications, in: Applied Chaos, eds Kim JH, Stringer J (Wiley, 1992) 3–30.

[FGS] Fiedler B, Gröger K, Sprekels J (eds), Equadiff'99, vol.2 (World Sci, 2000).

[Fr] Franks JM, Time dependent stable diffeomorphisms, Invent Math 24 (1974) 163–72.

[Gu] Gundlach M, Random homoclinic orbits, Random & Comput Dyn 3 (1995) 1–33.

[HV] Hale JK, Verduyn Lunel SM, Introduction to functional differential equations (Springer, 1993).

[J] Johnson RA, Ergodic theory and linear differential equations, J Diff Eq 28 (1978) 23–34.

[JN] Johnson R, Nerurkar M, Controllability, stabilization and the regulator problem for random differential systems, Mem Am Math Soc 136 (1998) 1–48.

[K] Kato T, Perturbation Theory for linear operators (Springer, 2nd ed, 1976)

[KH] Katok A, Hasselblatt B, Introduction to the modern theory of dynamical systems (Cambridge, 1995)

[LRS] Langa JA, Robinson JC, Suarez A, Stability, instability and bifurcation phenomena in non-autonomous differential equations, Nonlin 15 (2002) 1–17.

[MS] Massera JL, Schäffer JJ, Linear differential equations and functional analysis I & II, Ann Math 67 (1958) 517–73; 69 (1959) 535–74.

[MSe] Meyer KR, Sell GR, Melnikov transforms, Bernoulli bundles and almost periodic perturbations, Trans Am Math Soc 314 (1989) 63–105.

[RM] Rahn CD, Mote CD, On the stability of the damped Hill's equation with arbitrary bounded parametric excitation, Trans ASME 60 (1993) 366–

[RGO] Romeiras FJ, Grebogi C, Ott E, Multifractal properties of snapshot attractors of random maps, Phys Rev A 41 (1990) 784–99

[SS] Sacker RJ, Sell GR, Existence of dichotomies and invariant splittings for linear differential equations I,II&III, J Diff Eq 15 (1974) 429–58; 22 (1976) 478–96 and 497–522; and Sacker RJ, IV, J Diff Eq 27 (1978) 106–37.

[Se] Sell GR, Non-autonomous differential equations and dynamical systems, Am Math Soc 127 (1967) 241–83.

[Sie] Siegmund S, Dichotomy spectrum for non-autonomous differential equations, J Dyn Diff Eq 14 (2002) 243–

[Sim] Simiu E, Chaotic transitions in deterministic and stochastic dynamical systems (Princeton, 2002)

[Son] Sontag ED, Mathematical control theory (Springer, 1998)

[STUL] Stewart HB, Thompson JMT, Ueda Y, Lansbury AN, Optimal excape from potential wells – patterns of regular and chaotic bifurcation, Physica D 85 (1995) 259–95.

[TSU] Thompson JMT, Stewart HB, Ueda Y, Safe, explosive and dangerous bifurcations in dissipative dynamical systems, Phys Rev E 49 (1994) 1019–27.

[TRS] Thompson JMT, Rainey RCT, Soliman MS, Mechanics of ship capsize under direct and parametric wave excitation, Phil Trans R Soc Lond A 338 (1992) 471–90.

[UW] Uzunoglu V, White MH, The synchronous oscillator: a synchronization and tracking network, IEEE J Solid-state Circuits SC-20 (1985) 1214–25.

[W] Wiggins S, Normally hyperbolic invariant manifolds in dynamical systems (Springer, 1994)