Non-autonomous saddle-node bifurcations

URSS Summer project report

Gabriel Rémond-Tiedrez

Summer 2022

Contents

0 Introduction

This document is a report on my summer research project, done between the 3rd and 4th year of my undergraduate studies at Warwick, on non-autonomous saddle-node bifurcations.

It was supervised by Prof. Robert MacKay and funded by the URSS (Undergraduate Research Support Scheme), organised by the University of Warwick. Thank you to Robert for the support and answers to my many many questions. Thank you to the University of Warwick's URSS for giving me an opportunity to do some research.

This project changed directions a few times, when the problems we encountered didn't seem to budge. As such, I will attempt to recreate the evolution of the project as we went through it, noting the important snags along the way.

Note that this is by no means a research paper, simply a report of my summer of research. I therefore try to explain the most I can, to make this readable without any other references. However, I have given the reader a good number of references if more context is needed. If there are any questions or anything to add, feel free to contact me on gabriel.remond-tiedrez@warwick.ac.uk and I'd be very happy to have a chat.

1 A general functional analytic approach

This approach on hyperbolicity is taken from [Bis97] and [BM03]. For another explanation and some proofs, the reader is directed to the discussions on uniform hyperbolicity in both. Other useful references are the papers by Sacker and Sell [SS74, SS76a, SS76b, Sac78], but they consider a skew-product which we're avoiding for now.

1.1 Hyperbolic solutions

For autonomous systems $\dot{x} = f(x)$ with $f \in \mathcal{C}^1(\mathbb{R}^n, \mathbb{R}^n)$, the simplest solutions are equilibria x_0 with $f(x_0) = 0$. Looking at $\mathrm{D}f_{x_0}$ and its eigenvalues tells us whether it is hyperbolic, which is the case if and only if none of the eigenvalues are on the imaginary axis. We can then split the tangent space into 2 subspaces, E^- that contracts exponentially forwards in time and E^+ that contracts exponentially backwards in time. The former corresponds to eigenvalues with negative real part, the latter to eigenvalues with positive real part. Traditionally, E^- and E^+ are called the stable and unstable subspaces, respectively. But usually E^- is not stable and E^+ might be attracting, so we'll drop those names.

This equivalence of splitting and the condition on the eigenvalues only works for autonomous systems, however. A simple example of this is in [Cop78, p. 3]. This encourages finding another formulation of hyperbolicity, hopefully one that can be extended to non-autonomous systems.

Define $\mathcal{C}^0$ to be the space of bounded continuous functions from $\mathbb{R}$ to $\mathbb{R}^n$, equipped with the norm $\|x\|_{\mathcal{C}^0} = \sup_{t\in\mathbb{R}}|x(t)|$. Similarly, define $\mathcal{C}^1$ to be the space of continuously differentiable functions from $\mathbb{R}$ to $\mathbb{R}^n$, with

bounded norm and bounded derivative. For $x \in \mathcal{C}^1$, we have $\|x\|_{\mathcal{C}^1} = \max(\|x\|_{\mathcal{C}^0}, \|\dot{x}\|_{\mathcal{C}^0})$ and we want $\|x\|_{\mathcal{C}^1} < \infty$.

Theorem 1 (Theorem 1 in [BM03])**.** *An equilibrium x_0 is hyperbolic if and only if the linear operator $L : \mathcal{C}^1 \to \mathcal{C}^0, \xi \mapsto \dot{\xi} - \mathrm{D}f_{x_0}\xi$ is invertible.*

For any $\varphi \in \mathcal{C}^0$, we want to find a unique $\xi \in \mathcal{C}^1$ such that $\dot{\xi} - \mathrm{D}f_{x_0}\xi = \varphi$, or equivalently that ξ solves the inhomogeneous linear equation $\dot{\xi} = \mathrm{D}f_{x_0}\xi + \varphi$.

Note that we are only interested in bounded functions in both domains, so we are not solving an initial value/boundary problem, but rather finding the unique bounded solution.

Example 1 (Actually a non-example). Take the system

$$\begin{cases} \dot{x} = & -y \\ \dot{y} = & x \end{cases} \tag{1}$$

which has an equilibrium at $(0,0)$ and with $\mathrm{D}f_{(0,0)} = \left(\begin{smallmatrix} 0 & -1 \\ 1 & 0 \end{smallmatrix}\right)$, which has strictly imaginary eigenvalues $\pm i$. The origin is a centre, with no contraction in either direction of time. We consider the operator L as above and attempt to find the pre-image of $\varphi : \mathbb{R} \to \mathbb{R}^2, t \mapsto (0,0)$, solving $\dot{\xi} = \mathrm{D}f_{(0,0)}\xi$, the original equation. This has infinitely many bounded solutions so φ has no unique pre-image.

Now consider a non-autonomous dynamical system $\dot{x} = f(x,t)$ with $f : \mathbb{R}^n \times \mathbb{R} \to \mathbb{R}^n$, differentiable in x, with $\mathrm{D}f$ – the x-derivative – and f both continuous in (x,t).

Definition 1 (Definition 1 in [BM03])**.** A solution $x_0(t)$ of the above dynamical system is called *hyperbolic* if the linear operator $L : \mathcal{C}^1 \to \mathcal{C}^0, \xi \mapsto \dot{\xi} - \mathrm{D}f_{x_0(t),t}\xi$ is invertible.

It can be shown that this definition is equivalent to having a splitting of the tangent space into forwards and backwards contracting subspaces.

1.2 Saddle-node

Our interest is saddle-nodes, which are by definition not hyperbolic. They are non-hyperbolic in the simplest sense. In autonomous systems, we say that x_0 is a saddle-node for the 1D 1-parameter family of dynamical systems $\dot{x} = f(x,\lambda)$ at λ_0 if:

- x_0 is an equilibrium for the system with λ_0, $f(x_0,\lambda_0) = 0$,
- x_0 is not hyperbolic for the system with λ_0, the x-derivative $f_x(x_0,\lambda_0) = 0$,

- the 2^{nd} order x-derivative $f_{xx}(x_0, \lambda_0) \neq 0$,
- the λ-derivative $f_\lambda(x_0, \lambda_0) \neq 0$.

In higher dimensions, we would have a simple eigenvalue 0 of the Jacobian and no other eigenvalues on the imaginary axis. Then we would look at the dynamics on the 1D centre manifold, which would be as above. The extension is therefore not very complicated but it adds uninteresting baggage.

The reason we consider this to be the simplest non-hyperbolic behaviour is because, in the general higher-dimensional case, having a 1D centre manifold is the easiest non-hyperbolic option. Then in that 1 dimension, the easiest behaviour is when only need the first few orders of the Taylor expansion of the vector field, i.e. that the coefficients of the first order in λ and second order in the centre variable are non-zero. This is exactly what we check with the conditions on the derivatives.

The question now arises of how to express this in the functional analytic language we introduced. What happens with the operator L? It is certainly not invertible because x_0 is not hyperbolic, but we want it to be close.

We therefore focus on solutions $x_0(t)$ for which the operator $L : \mathcal{C}^1 \to \mathcal{C}^0, \xi \mapsto \dot{\xi} - \mathrm{D}f_{x_0,\lambda_0}\xi$ has a simple eigenvalue 0, or equivalently a 1D kernel.

Example 2. The canonical example of a saddle-node is the system $\dot{x} = x^2 + \lambda$, which has a saddle-node solution $x_0 \equiv 0$ at $\lambda_0 = 0$. The operator L here is just differentiation, whose kernel is the 1D subspace of constants. We could also consider $\dot{x} = x^3 + \lambda$, which would have the same operator L. However this is not a saddle-node as there is no second order term in x in the Taylor expansion.

We singled out it the property that the Jacobian at the equilibrium has a simple eigenvalue 0, and no others on the imaginary axis[1]. We can extend easily this to non-autonomous systems. It does not contain the information related to the form of the vector field on the centre manifold, so we will have to take care of that separately.

1.3 Implicit Function Theorem on a perturbation

Our approach is to consider a non-autonomous perturbation of an autonomous saddle-node. Consider the 1D system $\dot{x} = f(x, \lambda, t, \varepsilon)$, where for $\varepsilon = 0$ there is no time-dependency, and there is a saddle-node solution x_0 at λ_0. Formally speaking, the domain of f is $\mathbb{R} \times I \times \mathbb{R} \times J$, where I is a small interval around λ_0 and J is a small interval around 0.

Once we turn on ε and have it non-zero, the non-autonomous behaviour kicks in, which we view as a perturbation of the autonomous system, using some smoothness and continuity conditions on f. We then want to find

[1] If there was another, $\ker L$ would have a higher dimension.

conditions to ensure that we keep the saddle-node solution as we change ε. We will use the Implicit Function Theorem (IFT) on an appropriate map.

We assume that f is continuous and all the derivatives we consider are well-defined and continuous. Since we are working in one dimension, we will write f_x instead of $\mathrm{D}f$ for the x-derivative.

Consider the map $\Phi : \mathcal{C}^1 \times I \times \mathcal{C}^1 \times J \to \mathcal{C}^0 \times \mathcal{C}^0 \times \mathbb{R}$,

$$(x(\cdot), \lambda, v, \varepsilon) \mapsto \begin{bmatrix} \dot{x}(\cdot) - f(x(\cdot), \lambda, \cdot, \varepsilon) \\ \dot{v}(\cdot) - f_x(x(\cdot), \lambda, \cdot, \varepsilon) \cdot v(\cdot) \\ \alpha v(\cdot) - 1 \end{bmatrix}, \tag{2}$$

where $\alpha : \mathcal{C}^1 \to \mathbb{R}$ is some covector, we will take $\alpha v = v(0)$ for simplicity. This is to normalise v. The dependence on time is replaced by a dot to emphasise we're working with the whole function, not just its value at some time t. For solutions to the system, the first component will be 0. The second component being 0 is if $v \in \ker L$.

The starting point is an autonomous saddle-node at $\varepsilon = 0$. There, $\Phi(x_0, \lambda_0, 1, 0) = \begin{bmatrix} 0 \\ 0 \\ 0 \end{bmatrix}$. Note that $v_0 = 1$ because $\ker L$ is the space of bounded functions, then the normalising α puts it at 1.

The problem we now want to solve is finding zeroes of Φ, which we can do by looking at an initial zero – the autonomous saddle-node – and use the IFT to get a continuation as we change ε. We are looking for $(x_\varepsilon, \lambda_\varepsilon, v_\varepsilon)$ such that $\Phi(x_\varepsilon, \lambda_\varepsilon, v_\varepsilon, \varepsilon) = 0$, so we need to check the derivative with respect to x, λ and v is invertible.

We calculate that at $(x_0, \lambda_0, v_0, 0)$,

$$\mathrm{D}\Phi : (\delta x, \delta\lambda, \delta v) \mapsto \begin{bmatrix} \delta\dot{x} - f_x \cdot \delta x - f_\lambda \cdot \delta\lambda \\ \delta\dot{v} - f_{xx} \cdot \delta x \cdot v_0 - f_{x\lambda} \cdot \delta\lambda \cdot v_0 - f_x \cdot \delta v \\ \alpha\delta v \end{bmatrix}$$

for tangent vectors $\delta x, \delta\lambda$ and δv, and all derivatives of f appearing are evaluated at $(x_0, \lambda_0, v_0, 0)$[2]. For simplicity, we will write ξ, μ, w instead and, giving

$$\mathrm{D}\Phi : (\xi, \mu, w) \mapsto \begin{bmatrix} \dot{\xi} - f_x \cdot \xi - f_\lambda \cdot \mu \\ \dot{w} - f_{xx} \cdot \xi - f_{x\lambda} \cdot \mu - f_x \cdot w \\ \alpha w \end{bmatrix}. \tag{3}$$

Choose $g, h \in \mathcal{C}^0, c \in \mathbb{R}$, we want to find unique (ξ, μ, w) such that $\mathrm{D}\Phi(\xi, \mu, w) = (g, h, c)$. The equation above can be simplified further, using that $(x_0, \lambda_0, v_0, 0)$ corresponds to an autonomous saddle-node. In particular, $f_x = 0$ for all t, f_λ, f_{xx} and $f_{x\lambda}$ are constants with respect to t.

The first component gives the simple ODE $\dot{\xi} = g + f_\lambda \cdot \mu$, which yields

$$\xi(t) = \xi(0) + \int_0^t g(s)\mathrm{d}s + f_\lambda \mu \cdot t.$$

[2] $v_0 = 1$ so we will not bother writing it.

The second component gives $\dot{w} = f_{xx} \cdot \xi + f_{x\lambda} \cdot \mu + h$, which yields

$$w(t) = w(0) + f_{x\lambda}\mu \cdot t + \int_0^t \big[f_{xx} \cdot \xi(s) + h(s)\big] \mathrm{d}s.$$

The third component gives $w(0) = c$.

We still need to choose $\xi(0)$ and μ to have ξ and w be bounded. Focus on ξ first: we want $|\int_0^t g(s)\mathrm{d}s + f_\lambda \mu \cdot t| < M$ for some $M > 0$, for all t. One way to do this would be to use the average $\bar{g} = \lim_{t\to\infty} \frac{1}{t}\int_0^t g(s)\mathrm{d}s$ and have $\mu = -\bar{g}/f_\lambda$. The issues with this are that the average might be different as we go so $-\infty$ or $+\infty$, or either average might not even exist in the first place.

Another solution to have this bounded is to change the norm on $\mathcal{C}^0$ to, say, $\|g\|_{\mathcal{C}^0_n} = \sup_{t\in\mathbb{R}} |t|^{-n}|g(t)|$, to allow for functions to grow at most like $|t|^n, n \in \mathbb{N}$. But when integrating g, one gets at most $|t|^{n+1}$, which would grow too quickly to be bounded with the $\|\cdot\|_{\mathcal{C}^0_n}$ norm.

Similarly, one could change the norm to $\|g\|_{\mathcal{C}^0_\beta} = \sup_{t\in\mathbb{R}} e^{-\beta|t|}|g(t)|$, to allow growth of at most $e^{\beta|t|}, \beta > 0$. Then for $|g(t)| \leq Ce^{\beta|t|}$, we have $|\int_0^t g| \leq \int_0^t |g| \leq C\int_0^t e^{\beta|s|}\mathrm{d}s \leq \frac{C}{\beta}(e^{\beta|t|} - 1)$. We proved that if $\|g\|_{\mathcal{C}^0_\beta} < \infty$, then $\|\int_0^t g\|_{\mathcal{C}^0_\beta} < \infty$.

We can find constants C_1, C_2 such that $|g(t)| \leq C_1 e^{\beta|t|}$ and $|h(t)| \leq C_2 e^{\beta|t|}$. Then

$$|\xi(t)| \leq |\xi(0)| + \Big|\int_0^t g(s)\mathrm{d}s\Big| + |f_\lambda \mu| \cdot |t|,$$

so $\|\xi\|_{\mathcal{C}^0_\beta} < \infty$ from the comment above. We also have

$$|w(t)| \leq |w(0)| + |f_{x\lambda}\mu| \cdot |t| + |f_{xx}| \cdot \Big|\int_0^t \xi\Big| + \Big|\int_0^t h\Big|,$$

so $\|w\|_{\mathcal{C}^0_\beta} < \infty$. We therefore have both ξ, w bounded with the chosen norm.

The problem now comes from somewhere else: we are meant to find one unique pre-image that maps to (g, h, c) under $\mathrm{D}\Phi$, but instead we obtain a 2D subspace of them. Indeed, we get a bounded pre-image for any real numbers μ and $\xi(0)$. $\mathrm{D}\Phi$ is not invertible with this norm either.

1.4 Conclusions on this approach

There are possibilities of improvement or change in this method:

- One might change the norm to something else. We would need one with enough freedom to contain certain functions and their integrals, all the while maintaining enough restriction to obtain unique pre-images. Perhaps one could take a norm allowing functions to grow larger than polynomials but slower than exponentials.

- The way the problem was formulated here was to prepare it to the IFT, but one might want to use a Submersion Theorem argument or something else entirely.

- The type of solution we looked at might not have been nice enough to deal with. In particular, considering our infinite dimensional operator L with a simple eigenvalue at 0, a small perturbation might move a continuum of its eigenvalues around. Some might hit 0 and increase its multiplicity by an arbitrary amount. One would then be looking for conditions on the autonomous saddle-node solution for which this would not happen, i.e. 0 stays with multiplicity 1 when changing ε. In fact, our choice for what a saddle-node is might be wrong, or not restrictive enough.

Note that we start with an autonomous saddle-node at $\varepsilon = 0$. We do so because most derivatives we need to keep track of are just constants, making our analysis much easier.

If we do indeed get a continuation of the saddle-node solution, i.e. one with L having a simple eigenvalue 0, then f_{xx} and f_λ should be non-zero for all t by continuity and because they are non-zero constants in the $\varepsilon = 0$ case. Then we should get the same type of behaviour happening when varying λ, i.e. a saddle-node bifurcation.

The non-compactness of the direction of time seems to be a trouble-maker here, quickly causing unboundedness of solutions. We might therefore want to compactify time in some sense, making the following approach a natural progression.

One might want to try a similar approach with a discrete saddle-node instead of a continuous one, but it seems the problems encountered persist.

2 A discrete skew-product approach

This approach is taken from [AJ12], especially the use of their example on the 2-torus in the final section. The introductory part of this section is very similar to the aforementioned source, the reader is therefore encouraged to read it if interested. While the approach is entirely different, there is a discussion in [FJMV16] about the skew-product approach, which the reader might find worth it.

2.1 Skew-products

For this approach, we consider a 1-dimensional, discrete, skew-product dynamical system. We introduce a base space Θ and a forcing transformation $\omega : \Theta \to \Theta$, and study maps of the form

$$f : \Theta \times [a, b] \to \Theta \times [a, b], \ (\theta, x) \mapsto (\omega(\theta), f_\theta(x)). \tag{4}$$

The maps $f_\theta : [a,b] \to [a,b]$ are called *fibre maps*, where $[a,b]$ is our 1D phase space.

The objects studied are invariant graphs, i.e. functions[3] $\varphi : \Theta \to [a,b]$ such that $f_\theta(\varphi(\theta)) = \varphi(\omega(\theta))$. The authors then consider a 1-parameter family of maps $(f_\lambda)_{\lambda\in[0,1]}$ of the form (4) and a region $\Gamma \subset \Theta \times [a,b]$, and prove the existence of a critical bifurcation parameter λ_c such that:

- for $\lambda < \lambda_c$, f_λ has 2 invariant graphs in Γ;
- for $\lambda > \lambda_c$, f_λ has no invariant graphs in Γ;
- for $\lambda = \lambda_c$, f_λ either 1 or 2 invariant graphs in Γ. If there exist 2, then they are 'pinched'.

The precise statements and proofs are in Theorems 4.1 and 6.1 of [AJ12].

This is exactly analogous to the case of bifurcation of equilibria in the canonical saddle-node $\dot{x} = x^2 + \lambda$, where $\lambda_c = 0$:

- for $\lambda < \lambda_c$, there are 2 equilibria, one unstable, the other stable. This is similar to the 2 invariant graphs, one of which has an appropriate Lyapunov exponent positive, and the other invariant graph has a negative Lyapunov exponent;
- for $\lambda > \lambda_c$, there are no equilibria, similar to the absence of invariant graphs;
- for $\lambda = \lambda_c$, there is only one equilibrium, with strange behaviour. We could consider it as a merging of the previous two equilibria. Similarly, if we have 2 invariant graphs, they merge in some sense made rigorous by the definition of pinching.

The results of Theorems 4.1 and 6.1 are quite nice, giving the existence of a critical bifurcation value with straightforward assumptions. As mentioned in the end of sections 4 and 6, one could also look at ω-invariant subsets of Θ. For each such subset, we can get a different bifurcation value, and the critical value for the whole base space is effectively the infimum over invariant subsets.

2.2 Forcing with a cat map

We investigate the bifurcation values of orbits on the 2-torus under an Arnold cat map. An important difference between [AJ12] and what follows is that we do not study invariant graphs over the whole of the base space, but rather focus on specific invariant subsets, i.e. orbits.

[3] These functions should be measurable if Θ is a measure space, or continuous if Θ is a topological space.

Consider the base space $\Theta = \mathbb{T}^2 = (\mathbb{R}/\mathbb{Z})^2$ and the Arnold cat map ω represented by $\left[\begin{smallmatrix} 3 & 1 \\ 2 & 1 \end{smallmatrix}\right]$, that is, $\omega : \left[\begin{smallmatrix} x \\ y \end{smallmatrix}\right] \mapsto \left[\begin{smallmatrix} 3 & 1 \\ 2 & 1 \end{smallmatrix}\right]\left[\begin{smallmatrix} x \\ y \end{smallmatrix}\right] \mod 1$. We look at maps

$$f_{\lambda,\theta}(x) = \arctan(\alpha x) - \lambda - \gamma g(\theta),$$

where $\alpha > 1, \gamma > 0$ are constants[4] and $g : \mathbb{T}^2 \to \mathbb{R}$ is a smooth observable on the torus. We take $g(x, y) = \cos(2\pi x)\cos(2\pi y) + 1$.

There are really 2 saddle-node bifurcations with this map, one at the top, concave branch and the other at the bottom, convex branch of $\arctan(\alpha x)$. We will focus on the top one. An advantage of this map is that the bottom branch will be attracting throughout the parameter range of interest, meaning that solutions do not blow up in finite time. Even better, we know precisely where the trajectories go.

We can compute easily that ω has 2 fixed points on the torus: $(0, 0)$ and $(\frac{1}{2}, 0)$. The former is a global maximum of g and the latter a global minimum. The calculation of the bifurcation value on these fixed points is then straightforward, since this is then just an autonomous 1D saddle-node. We obtain $\lambda_c = \arctan(\sqrt{\alpha - 1}) - \gamma g - \frac{\sqrt{\alpha-1}}{\alpha}$, where g is to be evaluated at the fixed point. The earliest bifurcation is therefore the one at $(0, 0)$ and the latest one at $(\frac{1}{2}, 0)$[5].

To calculate the bifurcation value of other orbits under ω, we use numerical methods with MatLab. We work with periodic orbits, which coincide with rational points on the torus.

The initial method is by dichotomy, explained as follow. We start with an interval $[\lambda_0, \lambda_1]$ for the bifurcation parameter, for which we know that at λ_0, we still have both equilibria from the top branch and that at λ_1, both have cancelled each other out completely. We check behaviour at the midpoint $\frac{\lambda_0+\lambda_1}{2}$ of the interval: if there are no equilibria, we choose our new interval to be $[\lambda_0, \frac{\lambda_0+\lambda_1}{2}]$; and if the equilibria are still here, we choose $[\frac{\lambda_0+\lambda_1}{2}, \lambda_1]$. We repeat checking midpoints until we reach an acceptable level of accuracy. Our output is a small interval in which we know that the bifurcation happens, i.e. that we get the transition from 2 equilibria to none.

A few remarks on this:

- The choice of starting interval is important: if we choose one where no bifurcation happens, the process is obviously not going to yield a useful result. An easy option is to use the 2 fixed points and their seemingly extreme bifurcation values, to use them as starting endpoints. We could also take the range to be purposefully too large and add a few

[4] For discrete-time saddle-nodes, we need a point with derivative 1 and second derivative non-zero. We therefore need $\alpha > 1$: if $\alpha \in (0, 1)$, the derivative of $\arctan(\alpha x)$ is strictly bounded above by 1 and if $\alpha = 1$, the only place where the derivative is 1 has second derivative 0.

[5] It seems like these 2 bifurcation values are the extremes, but rigorously justify it seems beyond me.

iterations. This would not make the computations too much longer since the length of the interval decreases exponentially.

- To check existence of the equilibria, we do not find them explicitly but rather use the attracting properties of one equilibrium and of the bottom branch. We start with $x = 3$ in our phase space and iterate f a given number of times[6]. If the attracting equilibrium is still here, the trajectory will tend towards it, and will stay above or at 0. If the equilibria are gone, it will go towards the bottom branch, well below 0. Our check is simple, we look at whether we are above or below 0 after a number of iterations.

Figures (1a),(1b) and (1c) represent solutions in the phase space at different parameter values. The y-axis is the 1D phase space and the x-axis counts the number of iterations done with f.

2.3 A Newton step method

The method used above generally only uses a relatively small number of possible iterations of f. The interest of working with periodic orbits is also to approximate aperiodic ones, for which we want the periods to grow quite large. However, the efficiency of the method is low when looking at large periods, especially if we want to iterate enough for at least one period.

This encourages another approach, one using a Newton step method. We will make the saddle-node solution at the bifurcation value correspond to a zero of a smooth function Φ. Start at a seed $\mathbf{x}_0$, i.e. a point close to a zero of Φ and iterate the process defined by $\mathbf{x}_{n+1} = \mathbf{x}_n - (\mathrm{D}\Phi)^{-1}(\Phi(\mathbf{x}_n))$. If the derivative $\mathrm{D}\Phi$ is invertible and not too close to the $\mathbf{0}$ matrix, the $\mathbf{x}_n$'s converge to a nearby zero of Φ.

Consider an ω-orbit on the 2-torus of period $p \in \mathbb{N}$, starting at $\theta \in \Theta$. Now consider a sequence of p points $x_0, x_1, \ldots, x_{p-1} \in \mathbb{R}$, our phase space. We also need to pick a value for the parameter λ. We then define

[6] We choose $x = 3$ to ensure that we are not under both equilibria, or in between. Under both would immediately makes us go to the bottom branch, which is an issue. Between both would not be a problem, but the space between the equilibria will get very small, so we could easily fall under them if we are not careful. The solutions also contract very quickly so we do not lose much time from starting higher. We therefore go for the safer option and waste a little processing power.

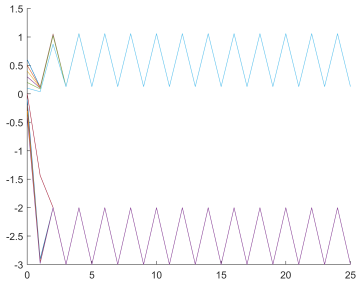

(a) $\lambda < \lambda_c$: the top attracting equilibrium is clearly visible. Solutions starting under the repelling one go towards the bottom branch.

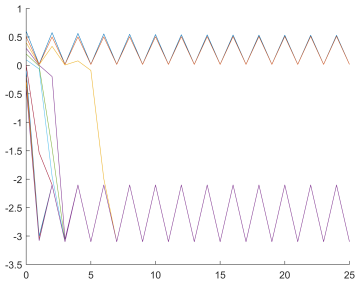

(b) $\lambda \approx \lambda_c$: the top saddle-node equilibrium is visible, note how trajectories are slower to go towards the bottom branch compared to the previous case.

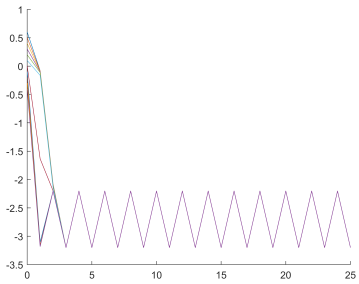

(c) $\lambda > \lambda_c$: both top equilibria have fully annihilated each other, causing all solutions to go towards the bottom branch.

Figure 1: Phase portraits for forcing with period 2 orbit starting at $(\frac{1}{2}, \frac{1}{2}) \in \Theta$, for $\alpha = 100$ and $\gamma = 0.5$. Note the periodic behaviour in all 3 figures.

$\Phi : \mathbb{R}^p \times \mathbb{R} \to \mathbb{R}^p \times \mathbb{R}$,

$$\Phi : (x_0, x_1, \ldots, x_{p-1}, \lambda) \mapsto \begin{bmatrix} f_{\lambda,\theta}(x_0) - x_1 \\ f_{\lambda,\omega(\theta)}(x_1) - x_2 \\ f_{\lambda,\omega^2(\theta)}(x_2) - x_3 \\ \vdots \\ f_{\lambda,\omega^{p-1}(\theta)}(x_{p-1}) - x_0 \\ \sum_{i=0}^{p-1} \log(f'_{\lambda,\omega^i(\theta)})(x_i) \end{bmatrix},$$

where f' is the derivative with respect to x.

A zero of this map means that $x_0, x_1, \ldots, x_{p-1}$ is a periodic orbit in the phase space for the value of λ. We obtain this from the first p components. The final component is to ensure this orbit is a saddle-node: a discrete-time saddle-node, say for a 1D system $y_{n+1} = h(y_n)$, is characterised by having first derivative $h' = 1$. The analogue of that for periodic orbits in the phase

space, is to have the derivative of all the compositions to be 1. The chain rule then yields a product of derivatives, evaluated at different points on the orbit. Taking log of both sides, we obtain that the sum of log's of derivatives is 0.

Computing the derivative $\mathrm{D}\Phi$ is straightforward:

$$\mathrm{D}\Phi = \begin{bmatrix} f'(x_0) & -1 & 0 & \dots & 0 & -1 \\ 0 & f'(x_1) & -1 & \dots & 0 & -1 \\ 0 & 0 & f'(x_2) & \dots & 0 & -1 \\ \vdots & \vdots & \vdots & \ddots & \vdots & \vdots \\ -1 & 0 & 0 & \dots & f'(x_{p-1}) & -1 \\ \frac{f''(x_0)}{f'(x_0)} & \frac{f''(x_1)}{f'(x_1)} & \frac{f''(x_2)}{f'(x_2)} & \dots & \frac{f''(x_{p-1})}{f'(x_{p-1})} & 0 \end{bmatrix},$$

where the dependence of f on λ, θ is suppressed because it disappears when differentiating w.r.t. x. We compute that $f'(x_i) = \frac{\alpha x_i}{1+\alpha^2 x_i^2}$ and $\frac{f''(x_i)}{f'(x_i)} = \frac{1-\alpha^2 x_i^2}{x_i(1+\alpha^2 x_i^2)}$.

We can now use the first method to get a basic approximation of a zero of Φ, then use the Newton step to get more accurate results. An important thing to note is that this method is unfortunately more unstable. For the parameters used in the phase portrait figures ($\alpha = 100$ and $\gamma = 0.5$), all orbits with period greater than 10 that were studied did not provide any usable results. The problem seemed to be that while most entries in the matrix are 0, some of them could get very large, which meant dealing with the matrix was complicated. For example, for $\alpha \gg 1$ and x_i close to 0, $\frac{f''(x_i)}{f'(x_i)}$ is large and not too uncommon.

A way to avoid this problem is to take smaller values, for example $\alpha = 10$ and $\gamma = 1 \times 10^{-3}$. All periodic orbits on Θ that were considered were manageable with this choice.

2.4 Eigenspaces of the origin

Now that we have a way to calculate accurately what the corresponding bifurcation value is, depending on the forcing orbit, we can navigate the rational points of the torus. Figures (2a) and (2b) show how different starting points on the torus give different bifurcation values. There are a couple of things to point out.

Focus on the point $(0, 0)$ in (2a). Note that this is a flat representation of the torus, so the 4 corners $(0, 0)$, $(1, 0)$, $(0, 1)$ and $(1, 1)$ are really one and the same. One can calculate eigenvectors of ω to be $\begin{bmatrix} 1 \\ \sqrt{3}-1 \end{bmatrix}$ and $\begin{bmatrix} (1-\sqrt{3})/2 \\ 1 \end{bmatrix}$, where the former corresponds to the backwards contracting direction of ω and the latter the forwards contracting direction. We can then notice that these directions are exactly the directions of the red streaks emanating from $(0, 0)$ and all other corners. We get the same behaviour at the other fixed point $(\frac{1}{2}, 0)$ and $(\frac{1}{2}, 1)$, but with blue streaks instead.

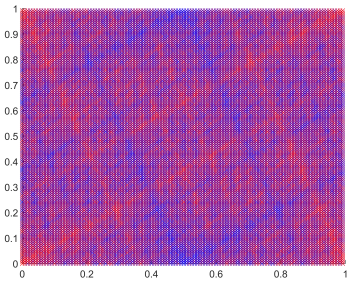

(a) Rational points with denominator 120.

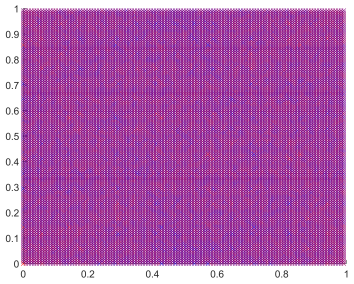

(b) Rational points with denominator 119.

Figure 2: Bifurcation values for different starting $\theta \in \Theta$, red means the bifurcation happens early, blue means late.

This seems to indicate that being on the eigenspace of a fixed point makes the bifurcation occur at the same time. We could imagine that once $(0, 0)$ bifurcates, for example, both its eigenspaces do so with it. It would then leave only a Cantor set of points on the torus, which bifurcate later.

This image is quite an elegant one but the second figure tells another story. There seems to be less obvious structure in (2b). The picture is a lot messier, and most points are mixed between red and blue. So what happened here? This is still unclear, but a possibility is that this is to do with the period of the points we considered. Indeed, rational points with denominator 120 have period at most 12, whereas they can reach periods of 72 for denominator 119. However, the reason for which this matters is still unsure.

We investigate the eigenspaces of the origin in one other way, by taking rational points converging to a point on an eigenspace, and calculating their corresponding bifurcation values. A rational point very close to the eigenspace should have quite a large period and stay close to the point on the eigenspace for an good number of iterations, therefore have a similar bifurcation value. More precisely, we have rational points (x_n, y_n) that converge to a point (x_∞, y_∞) on the eigenspace. We then look at the bifurcation values for each (x_n, y_n) and consider the bifurcation value of (x_∞, y_∞) to be the limit of that sequence.

Consider the function $\lambda_c : \mathbb{T}^2 \to \mathbb{R}$, taking a point (x, y) on the torus and whose output is the corresponding bifurcation value. The above is equivalent to asking that λ_c is continuous at (x_∞, y_∞), a risky assumption.

There are different ways to converge to a point on an eigenspace – be it the backwards or forwards contracting one. We fix a rational value x_∞ of x and choose the corresponding value y_∞ of y, to have (x_∞, y_∞) on the eigenspace of interest. We then choose our sequence to be $(x_n, y_n) =$

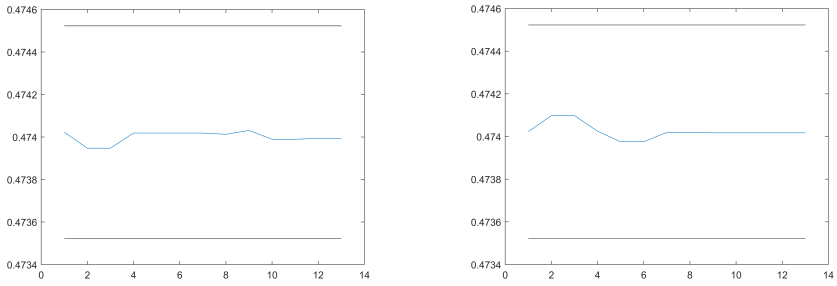

(a) Forward contracting eigenspace, with $x_\infty = 0.3$.

(b) Backwards contracting eigenspace, with $x_\infty = 0.2$.

Figure 3: Bifurcation values for points approaching eigenspaces of the origin. The x-axis is the n of our sequence and the y-axis is the corresponding bifurcation value. The 2 flat lines above and below are the bifurcation values of $(\frac{1}{2}, 0)$ and the origin, respectively. These are plotted since they are bounds for the possible bifurcation values. We therefore get a better idea of 'when' the bifurcation happens.

$(x_\infty, \frac{\lfloor 2^n y_\infty \rfloor + 1}{2^n})$[7], all of which are rational points quickly converging to (x_∞, y_∞).

Figures (3a) and (3b) show something other than what was expected[8]. Indeed, if the eigenspaces bifurcate at the same time as the origin and that λ_c is continuous at (x_∞, y_∞), we would expect the curve to get close to the bottom line, that is to tend towards the bifurcation value of $(0, 0)$. It seems that at least one of these assumptions is wrong. Ideally, we would calculate the bifurcation value of points on the eigenspaces directly, but it is numerically impossible because of their aperiodicity. We therefore do not have a way of checking that condition. The curves obtained seem instead to go towards some average value.

An interesting aspect to note is that as we go further in the approximations, the period increases. In general, multiplying the denominator at least increases the period. In contrast to that, the size of the denominator itself is not directly correlated with the period, as seen in figure (2). Consider the function $\lambda_c : \mathbb{R}^p \to \mathbb{R}$, taking p values of periodic forcing and outputting the corresponding bifurcation value. This function would depend crucially on p, seemingly giving values close to the average as p increases. Evidence for that is the increased proportion of purple in figure (2b) and the value that the curves of figure (3) seem to go towards.

There are no proofs here, only observations and a few ideas. To make

[7]The $+1$ in the numerator is not necessary, but it was used to obtain the figures so it is written here for coherence.

[8]Other points on the eigenspaces were considered as well, but they all yielded a similar picture, so are not shown here.

the problem simpler, we take another function f, which we could consider as a modified limit when taking the parameter α to infinity.

2.5 Saddle-node on a step function

For $\alpha \gg 1$, the function $\arctan(\alpha x)$ is quite close to a step-function from $-\frac{\pi}{2}$ to $\frac{\pi}{2}$. The factor of $\frac{\pi}{2}$ adds nothing interesting, so we consider a step function from -1 to $+1$:

$$\sigma(x) = \begin{cases} -1, & \text{if } x < 0, \\ +1, & \text{if } x \geq 0, \end{cases}$$

and consider the difference equation

$$x_{n+1} = \sigma(x_n) - \lambda - g_n. \tag{5}$$

Note that we dropped γ for simplicity, and because the analysis here won't need the forcing to be as weak as before. The forcing also changes form, we now work with a sequence $(g_n)_{n\in\mathbb{N}}$ of values of forcing. This could still be obtained with a skew-product but we do not use the structure provided by that.

An important choice for our step function is where we map 0. We could take -1, 0 or +1. Taking 0 would make σ a pointwise limit of $\frac{2}{\pi}\arctan(\alpha x)$ as $\alpha \to \infty$. However, we will choose +1. Thinking back to the $\alpha < \infty$ case with no non-autonomous forcing, a saddle-node point is the tangential intersection of the diagonal $y = x$ and $y = \frac{2}{\pi}\arctan(\alpha x) - \lambda$. We can calculate the x value for that quite easily using that it has derivative 1, giving $x_\pm = \pm\frac{2\sqrt{\alpha-1}}{\pi\alpha}$. The sign depends on which side of $\frac{2}{\pi}\arctan(\alpha x)$ we consider, so we consider x_+. Note that $\left(x_+, \frac{2}{\pi}\arctan(\alpha x_+)\right) \to (0, 1)$ as α goes to infinity, so choosing $\sigma(0) = +1$ is coherent. If we chose 0, we would not have any saddle-node behaviour, and if we chose -1, we would be working with the bottom branch instead.

We now redefine what a saddle-node is for this case, as the usual characterisation using derivatives needs some adjusting. We still ignore the forcing for now. The point that would have derivative 1 is at $x = 0 = \lim_{\alpha\to\infty} \frac{\sqrt{\alpha-1}}{\alpha}$, and the required intersection with $y = x$ means that one of the branches of $\sigma(x) - \lambda$ must hit 0. Without any forcing, we would have $\lambda = \pm 1$, depending on which branch hits 0.

Now for the case with forcing: we say that a sequence $(x_n)_{n\in\mathbb{N}}$ is a saddle-node solution for aperiodic forcing $(g_n)_{n\in\mathbb{N}}$ if $x_n = 0$ for exactly one n; similarly for periodic forcing $(g_n)_{n\in\mathbb{N}}$, we have a saddle-node if $x_n = 0$ exactly once per period. One might have saddle-nodes with more than one $x_n = 0$ if multiple values of forcing are equal, but this case is not different enough to warrant its own discussion.

The analysis is now straightforward and we obtain simple conditions on the forcing for which we have saddle-nodes, with a very simple way of obtaining the critical value for λ.

2.6 Using a simple lemma

We first consider the case of periodic forcing $(g_n)_{n\in\mathbb{N}}$ of period p, with $g_n \in (0,2)$ for all n. That is, $g_{n+p} = g_n$ for all n and p is the smallest such integer[9]. We assume that we already have a saddle-node solution, then investigate what properties the x_n's and g_n's have to satisfy.

Consider a saddle-node solution $(x_n)_{n\in\mathbb{N}}$ of period p, for which we choose to have $x_p = 0$ and $x_n \neq 0$ for all $n \neq p$. We first prove a small lemma, which will simplify the approach.

Lemma. *Let $n \in \mathbb{N}$, then we must have either*

$$\sigma(x_n)\cdot\sigma(x_{p-1}) \geq 0 \text{ or } \sigma(x_n)\cdot x_{n+1} \geq 0. \tag{6}$$

Proof. We proceed by contradiction: suppose that there exists a $n \in \mathbb{N}$[10] with $\sigma(x_n)\cdot\sigma(x_{p-1}) < 0$ and $\sigma(x_n)\cdot x_{n+1} < 0$. Consider the 2 possibilities for $\sigma(x_j)$:

- $\sigma(x_n) = 1$: then we deduce that $\sigma(x_{p-1}) = -1$ and that $x_{n+1} < 0$. We calculate that $-\lambda = g_{p-1} - \sigma(x_{p-1}) = g_{p-1} + 1$, then see that $x_{n+1} = 2 + g_{p-1} - g_n$. Using the sign yields $g_n - g_{p-1} > 2$, which is a contradiction to our bounds on the forcing.
- $\sigma(x_n) = -1$: then $\sigma(x_{p-1}) = 1$ and $x_{n+1} > 0$. Here, $-\lambda = g_{p-1} - 1$, so $x_{n+1} = -2 + g_{p-1} - g_n$. Using the sign yields $g_{p-1} - g_n > 2$, which is a contradiction.

□

This result is very useful for what follows, but we will adjust it. Since σ takes only 2 values, we can change the first condition into $\sigma(x_n) = \sigma(x_{p-1})$. Then since this is obviously true for $n = p-1$, we can assume that $n \neq p-1$ to simplify the second condition into $\sigma(x_n)\cdot x_{n+1} > 0$. We're using here that x_p is the only point in the period that hits 0[11]. So the final condition is: for all $n \in \mathbb{N}$, we have either

$$\sigma(x_n) = \sigma(x_{p-1}) \text{ or } \sigma(x_n)\cdot x_{n+1} > 0. \tag{7}$$

[9] In the skew-product image, the g_n's would be the values of an observable evaluated at points of a periodic orbit on Θ.

[10] For this periodic forcing, we only really need to consider p values, but choosing any n makes it clear that this lemma works for aperiodic forcing as well.

[11] In the case of aperiodic forcing, we would be using that x_p is the only point that hits 0.

Calculating the bifurcation value is very easy. If $x_p = 0$, (5) yields $\lambda_c = \sigma(x_{p-1}) - g_{p-1}$. We could therefore expect g_{p-1} to be some special value of forcing. Another useful fact is that $\sigma(x_p) = \sigma(0) = 1$, which we can use alongside condition (7) to check different possible signs the x_n's can take. Some choices will be compatible with each other, others will give contradictions.

We split the analysis in 2 cases:

- $\sigma(x_n) = \sigma(x_{p-1})$ for all n: in particular $\sigma(x_p) = \sigma(x_{p-1}) = 1$, so we have $x_n > 0$ for $n = 1, 2, \ldots, p-1$.
- there exist n such that $\sigma(x_n) \neq \sigma(x_{p-1})$: then (7) implies that $\sigma(x_n) \cdot x_{n+1} > 0$, which in turn means that $\sigma(x_{n+1}) \neq \sigma(x_{p-1})$. So similarly $\sigma(x_{n+2}) \neq \sigma(x_{p-1})$ and so on. We get a chain of x_n's all with a sign different from x_{p-1}. Since we have periodic forcing, this means that this chain will loop around all the way up to x_{p-2}, independently of where the first different sign appears. But this yields a contradiction to (7), because $\sigma(x_{p-2}) \neq \sigma(x_{p-1})$ and $\sigma(x_{p-2}) \cdot x_{p-2+1} < 0$. This case will split into three when looking at the aperiodic case, and we will need to look at the first occurrence of the different sign.

With periodic forcing, only one possibility is valid, that of $x_n > 0$ for all $n \neq p$. We calculate $\lambda_c = \sigma(x_{p-1}) - g_{p-1} = 1 - g_{p-1}$. Then for any $n \neq p$, we have $0 < x_n = \sigma(x_{n-1}) - \lambda_c - g_{n-1} = 1 - 1 + g_{p-1} - g_{n-1}$, i.e. we require $g_{p-1} > g_{n-1}$ for all $n \neq p$.

So for given periodic forcing g_n, the corresponding saddle-node solution must have $x_n \geq 0$ for all n, and has $x_p = 0$ if and only if the value g_{p-1} of forcing is strictly maximal. We haven't quite proven the equivalence, but it is straightforward. Note that the condition of g_{p-1} being strictly larger than any other value can be relaxed: if we allow equality, we could get multiple x_n's to hit 0. In fact we have a 0 after each maximal value.

A final remark on this periodic case: the condition that $x_n \geq 0$ for all n means that we stay on the top branch of our step function. The fact that we need to start on it should be comforting news, since it is the branch we were focussing on with the $\arctan(\alpha x)$ example. But starting on it means we must periodically return to the top branch, which then requires the solution to stay on it for all n. In a completely symmetric argument, if we chose to map 0 to -1 under the step function, we would need to stay on the bottom branch.

2.7 Extending to aperiodic forcing

Most of the work was done in the lemma giving (7). We introduce a slight shift in notation, calling the index where x_n hits 0 N instead of p. That is, $x_N = 0$ and $x_n \neq 0$ for $n \neq N$. Just like in the periodic case, we split into 2 cases, depending on if all the signs of $\sigma(x_n)$ match or not:

- $\sigma(x_n) = \sigma(x_{N-1})$ for all n: in particular $\sigma(x_N) = \sigma(x_{N-1}) = 1$, so we have $x_n > 0$ for $n = 1, 2, \ldots, N-1$. This is exactly as in the periodic case.

- there exist n such that $\sigma(x_n) \neq \sigma(x_{N-1})$: let m be the smallest such integer. Then (7) implies that $\sigma(x_m) \cdot x_{m+1} > 0$, which in turn means that $\sigma(x_{m+1}) \neq \sigma(x_{N-1})$. So similarly $\sigma(x_{m+2}) \neq \sigma(x_{N-1})$ and so on. Like in the periodic case, we get a chain of x_n's all of a sign different from x_{N-1}. The absence of periodicity means we cannot discard this possibility. We therefore require introducing m, a pointless addition in the periodic case. We divide this case further:

  - $m < N-1$: the chain of x_n's starting at x_m means that x_{N-2} has the opposite sign to that of x_{N-1}, i.e. $\sigma(x_{N-2}) \neq \sigma(x_{N-1})$. Then (7) implies that we have $\sigma(x_{N-2}) \cdot x_{N-2+1} = \sigma(x_{N-2}) \cdot x_{N-1} > 0$, yielding $\sigma(x_{N-2}) = \sigma(x_{N-1})$, a contradiction.
  - $m = N$: $\sigma(x_N) = 1$ so we must have $\sigma(x_{N-1}) = -1$. Thus the chain starting at $x_m = x_N$ must contain only positive x_n's, and $x_n < 0$ for $n < N$.
  - $m \geq N+1$: $\sigma(x_N) = 1$ and $\sigma(x_{N-1}) = \sigma(x_N)$, so for $j < m$ and $j \neq N$, $x_j > 0$. The rest is the chain emanating from x_m, all of which are negative.

So we have 3 possible cases:

- $x_j > 0$ for all $j \neq N$,
- $x_j < 0$ for all $j \leq N-1$ and $x_j > 0$ for all $j \geq N+1$,
- $x_j > 0$ for all $j < m, j \neq N$ and $x_j < 0$ for $j \geq m$.

I find it clearer written down in the following symbolic way:

- $+ \cdots + 0 + \cdots$,
- $- \cdots - 0 + \cdots$,
- $+ \cdots + 0 + \cdots + - \cdots$,

where each sign is the sign of the corresponding x_n.

Let us analyse what conditions on g_n are needed for each case.

2.7.1 $+ \cdots + 0 + \cdots$

We calculate $-\lambda_c = g_{N-1} - 1$.

Picking any $j \neq N-1$, we have $\sigma(x_j) = 1$ and $x_{j+1} > 0$. Then $x_{j+1} = g_{N-1} - g_j$ implies that $g_{N-1} > g_j$ for all $j \neq N-1$. This is exactly as in the periodic case, we need that g_{N-1} is the strict maximum.

2.7.2 $- \cdots - 0 + \cdots$

We calculate $-\lambda_c = g_{N-1} + 1$.

For $0 \leq j \leq N-2$, $\sigma(x_j) = -1$ and $x_{j+1} < 0$. Then $x_{j+1} = g_{N-1} - g_j$, so we need $g_{N-1} < g_j$.

For $j \geq N$, $\sigma(x_j) = 1$ and $x_{j+1} > 0$. Then $x_{j+1} = 2 + g_{N-1} - g_j$, so we need $g_j - g_{N-1} < 2$, which is fine since we chose our forcing in $(0, 2)$.

2.7.3 $+ \cdots + 0 + \cdots + - \cdots$

We calculate $-\lambda_c = g_{N-1} - 1$.

For $0 \leq j \leq N-2$, $\sigma(x_j) = 1$ and $x_{j+1} > 0$. Then $x_{j+1} = g_{N-1} - g_j$, so we need $g_{N-1} > g_j$.

If $m \geq N+2$, we get the following case for $N \leq j \leq m-2$, $\sigma(x_j) = 1$ and $x_{j+1} > 0$. Here we also need $g_{N-1} > g_j$.

For $j = m-1$, $\sigma(x_j) = 1$ and $x_{j+1} = x_m < 0$. Then $x_{j+1} = g_{N-1} - g_j$, so we need $g_{N-1} < g_j = g_{m-1}$.

For $j \geq m$, $\sigma(x_j) = -1$ and $x_{j+1} < 0$. Then $x_{j+1} = -2 + g_{N-1} - g_j$, so we need $g_{N-1} - g_j < 2$, which is fine because of our bounds on the forcing.

2.8 Another viewpoint

For this section, we will look at the forcing values first, then determine if we get a bifurcation and where the sequence x_n hits 0. We use exactly the same results but take another point of view.

Suppose that we have a g_{N-1} that is a strict maximum of all values of the forcing. Then we have the case in 2.7.1: $x_N = 0$, all the other x_n's are strictly positive and the critical value of the parameter is $\beta_c = 1 - g_{N-1}$.

Now suppose that we have a g_{N-1} that is a strict maximum of all values of forcing from g_0 to g_{m-2} for some $m - 2 > N - 1$, and that $g_{m-1} > g_{N-1}$. Then we are in the case of 2.7.3: $x_N = 0$, all n's for which $g_{N-1} > g_{n-1}$ have x_n positive, all the others are negative. Note here that all the values of forcing from m onwards do not matter, so long as $g_n \in (0, 2)$. The critical value of the parameter is $\beta_c = 1 - g_{N-1}$ as above.

One can think of this case as the more common version of the first, we still need a certain value of forcing to be maximal, but it can get "overthrown" and cause the x_n's to fall down on the bottom branch. The first case seems infinitely more unstable, as only one value of forcing larger than g_{N-1} can topple the result into the second case.

Lastly, suppose that we have a g_{N-1} that is the strict minimum of values of the forcing from g_0 to itself, with g_N. Then we are in the case of 2.7.2: $x_N = 0$, all n's before that are negative, all the others are positive.

At first sight, this case just seems slightly different to the previous ones, but it is more troublesome than one might think. The reason for that is uniqueness of solutions. In the first 2 cases, we get a unique combination

of signs that gives a saddle-node solution[12]. For this case, let us consider a sequence of forcing (g_n) that strictly decreases. Then one could start with negative x_0, and choose any n to have $x_n = 0$. Each one would give another entirely valid solution, all of whose bifurcation value is different.

We could also imagine a sequence strictly decreasing apart from some g_{N-1} being the strict maximum, with $N-1 > 0$. Then we could have $x_n = 0$ before x_N, because g_{n-1} would be the strict minimum of the values up to that point. In this case we'd start off with negative values for x_n. But we could also have $x_N = 0$ and we would start with positive values for x_n. The bifurcation value would also change considerably.

How does one justify discarding this problematic case? Note that we wanted to focus on solutions starting on the top branch, which is why we map 0 to +1 under the step function. But this case occurs precisely when we start on the bottom branch. If we fix that $x_0 \geq 0$, we are done and saddle-node solutions are uniquely determined by the sequence of forcing.

2.9 Conclusions on this approach

One question the reader could ask is why we require $g_n \in (0, 2)$ for all n. The condition we really need is that the maximal difference between any 2 values of forcing is strictly less than 2, in order to get the contradictions used. Then the choice of range is arbitrary, so we take a simple one. The reason that 2 is our bound is because of the step function considered. Looking back at the contradiction arguments in the lemma, one can notice that 2 is obtained as the difference between the 2 steps of σ. The intuition behind this is that if $|g_n - g_{n+1}| \geq 2$ for some n, the solution can jump from one branch to the other.

One might notice that this is exactly what happens in the latter case of aperiodic forcing, where the solution starts on the top branch, and jumps down once an appropriate value of forcing is reached. However, this is not caused by a large change in forcing, so is allowed. Indeed, we only need to have $g_{m-1} > g_{N-1}$, not a condition on their difference being large. When forcing with g_{N-1}, we hit 0 exactly, which by the choice of step function, brings us back to the top branch. When taking a larger value g_{m-1}, we go past 0 into negative reals, get mapped to the bottom branch and stay there.

Note that this argument again uses $\sigma(0) = +1$, an important choice.

Let us go back to the starting point of this section on skew-product systems, more specifically the Arnold cat map forcing used. What is done above was, as previously said, only observations and a few ideas. The main idea was to use the dynamics given by the Arnold cat map, mostly focussing on the eigenspaces, to determine information about bifurcation values in the phase space. More specifically, we computed the bifurcation values of

[12] This hasn't been formally done, but seems straightforward.

points near the eigenspaces of the origin, hoping they would be close to the bifurcation value of 0. The reality seems quite different, with the period of the starting point on the torus having an unsurprisingly important role.

An interesting follow up to this would be to more thoroughly investigate bifurcations over every orbit, or more generally invariant subsets. One could keep going with the example of a cat map, or with more general forcing transformations on a general base space, to fit into the framework of [AJ12].

3 Conclusion

We started with a general functional analytic approach to consider saddle-node solutions to autonomous systems, checking whether they remained under small time-dependent forcing. The unbounded direction of time caused problems, however. This is by contrast to the skew-product approach studied in [AJ12], where one has a lot more control over the time-dependence, writing it as a dynamical system on a base space. In particular, the examples worked through are both on compact spaces, where the direction of time cannot blow up.

The last part of this report looks at a special case of non-autonomous forcing, where one works with a step function. This is symbolic dynamics, with which the whole behaviour of the system can be understood. An interesting aspect to note is that while the forcing did not need to have any specific structure, we required specific bounds. However, the gap from differentiable systems to step functions might need some filling in.

References

[AJ12] V. Anagnostopoulou and T. Jäger. Nonautonomous saddle-node bifurcations: Random and deterministic forcing. *Journal of Differential Equations*, 253(2):379–399, 2012.

[Bis97] Zahir Bishnani. Safety criteria for aperiodic dynamical systems. July 1997.

[BM03] Z. Bishnani and R. S. MacKay. Safety criteria for aperiodically forced systems. *Dynamical Systems*, 18(2):107–129, 2003.

[Cop78] W. A. Coppel. *Dichotomies in Stability Theory*. Springer Berlin Heidelberg, Berlin, Heidelberg, 1978.

[FJMV16] Matteo Franca, Russell Johnson, and Victor Muñoz-Villarragut. On the nonautonomous Hopf bifurcation problem. *Discrete and Continuous Dynamical Systems - S*, 9(4):1119–1148, 2016.

[Sac78] Robert J. Sacker. Existence of dichotomies and invariant splittings for linear differential systems IV. *Journal of Differential Equations*, 27(1):106–137, 1978.

[SS74] Robert J. Sacker and George R. Sell. Existence of dichotomies and invariant splittings for linear differential systems I. *Journal of Differential Equations*, 15(3):429–458, 1974.

[SS76a] Robert J. Sacker and George R. Sell. Existence of dichotomies and invariant splittings for linear differential systems, II. *Journal of Differential Equations*, 22(2):478–496, 1976.

[SS76b] Robert J. Sacker and George R. Sell. Existence of dichotomies and invariant splittings for linear differential systems, III. *Journal of Differential Equations*, 22(2):497–522, 1976.