



Universiteit  
Leiden  
The Netherlands

## **Pulses in singularly perturbed reaction-diffusion systems**

Veerman, F.W.J.

### **Citation**

Veerman, F. W. J. (2013, September 25). *Pulses in singularly perturbed reaction-diffusion systems*. Retrieved from <https://hdl.handle.net/1887/21788>

Version: Not Applicable (or Unknown)

License: [Leiden University Non-exclusive license](#)

Downloaded from: <https://hdl.handle.net/1887/21788>

**Note:** To cite this publication please use the final published version (if applicable).

Cover Page



Universiteit Leiden



The handle <http://hdl.handle.net/1887/21788> holds various files of this Leiden University dissertation.

**Author:** Veerman, Frits

**Title:** Pulses in singularly perturbed reaction-diffusion systems

**Issue Date:** 2013-09-25

# 1

## Introduction and Summary

---

This introduction is written for non-scientists; its aim is to present the subject and research in this thesis in a way which is understandable for a broad audience. If you are a scientist, or even a mathematician, you might find this introduction lacking a certain depth. In that case, I recommend reading the introductory sections of the subsequent chapters – their content is specifically aimed at scientists, mathematicians in particular.

### 1.1 How should I read this?

There are a couple of reasons why a thesis in mathematics is hard to read for non-mathematicians. First, of course, there are the formulas. A mathematician conveys a large part of his message using symbols, and the connections between them by formulas. If you haven't got that much experience in using symbols and reading formulas, texts which heavily rely on them are notoriously hard, or even impossible, to read. Having said this, the use of symbols to convey mathematics is not only convenient, but also necessary. It allows the researcher to represent certain ideas (which might be quite abstract) using a few symbols, thereby keeping his or her reasoning accessible and clear to follow – for the fellow mathematician, of course. If you would try to reconvert all the symbols in this thesis into words, the text would very quickly spiral into incomprehensibility: the sentences would span several pages, it would be impossible to structure them in a clear way, giving up any hope for conveying the ideas you want your reader to understand. The conciseness and clarity of symbols, and their use in mathematics, have proved beneficial for many, many centuries. Symbols

and formulas allow you to discover new connections and relations, which opens the door to abstract thinking and deeper understanding of the subject you're considering – and that's not only true for mathematicians, but also for other scientists who use the 'language of mathematics' to summarise their findings.

When you start to work with symbols and formulas, and thereby gain some experience in reading them, you'll notice that these symbols (and the ideas they represent) become more tangible. More and more, you get an idea of what that symbol stands for. You get a feeling of how a symbol behaves, how it reacts to other symbols, what it does. Then, you can start to shove them around, manipulate them, and introduce new symbols because that's the best way to explain what you found – and suddenly, you're doing mathematics.

The second reason why mathematics is hard to read, is the language that is being used. Your goal as a mathematician is to convey objective truths, to tell a coherent and logically sound story. That means the language becomes objective as well: there is no 'I' or 'you' in mathematics, you can hope for a 'we' at best. In a mathematical text, you're taken along a route towards understanding a mathematical topic, guided by the author. Anything which might reek like subjectivity is to be avoided at all costs, is the opinion of many. Mathematical truth does (or should) after all not depend on who's presenting it. Phrasing your sentences in subjective form also makes you more vulnerable to criticism: *you* might say it is so, but that doesn't mean *I* have to believe it.

Although often deemed necessary, this practice doesn't do the readability of the mathematical text any good. As you might have already noticed, I've chosen a different style for this introduction. At a risk of being 'not scientific enough', i.e. not objective enough, I think it is necessary, if you want your ideas to be understood by a wider audience, to present ideas through a text which is accessible to the non-mathematical, non-scientific reader. That is exactly what I try to accomplish with this introduction. Sometimes the nature of the subject I'm describing is such that using an objective style is unavoidable; however, I'll try to refrain from doing so. Once the 'real' content starts in chapter 2, you'll notice a change in style from the somewhat direct, subjective style wielded in this introductory chapter, to the objective and somewhat indirect 'mathematical' style. As argued above, this is a necessary feature of mathematical research texts.

Text isn't everything. As a mathematician, I've noticed that deeper understanding of a phenomenon through symbol manipulation goes hand in hand with the development of a certain mental picture. Since the objects you're working with are often of an abstract nature, this mental picture cannot be more than approximately accurate.

In that respect, I consider myself lucky that I'm an applied mathematician. More often than in other, more pure branches of mathematics, I've got the possibility to accurately visualise the objects I'm analysing. Since the excessive use of symbols holds the danger of obscuring the analysis, the use of pictures can be beneficiary to understanding the topic at hand. One of my goals is therefore to give the reader some idea what the pictures in this thesis mean. If you're browsing through the mathematical chapters, you encounter a figure and think 'Ah! I've seen something like that before, could this be related?', that goal is achieved.

This introduction is structured as follows. Based on the consideration 'if you've got a thesis in your hand, you should at least be able to understand its title', I will first explain some concepts which are central to the research area in which the topic of this thesis falls (section 1.2). Along the way, all the words which make up the title will be introduced. This part is specifically written for non-mathematicians, even for non-scientists. As you can see while skimming through the introduction, there aren't as many formulas as you would expect from a mathematics thesis – especially in comparison with the next chapters, where the 'real' content can be found. Once I've told enough to explain the title, it's time to dive a little more into the research itself. If you've become interested at that point – and I hope you are –, I'll explain what the research presented in this thesis entails, since that is, purely based on the title, not at all clear. I'll explain in general the research methods and ideas used in this thesis (section 1.3). Also, I'll shed light on some research results, and tell you why they are important and why they are new. Although this part will unavoidably be a little bit more technical, it's still possible to get the message across without going into too much detail. In the end (section 1.4), I'll summarize the content of each chapter, giving you an overview of the content of this thesis.

## 1.2 Concepts

### 1.2.1 Patterns

What is a pattern? In the broadest sense, you could characterise a pattern as an 'observable regularity'. In nature, patterns are all over the place. The most obvious ones are stripes or spots on animal skins as found on zebras, leopards, cats, boar piglets; more elaborate spiral patterns occur as fingerprints, sea shells or snail shells. Once you start looking for patterns, 'there's something, then there's nothing, then there's

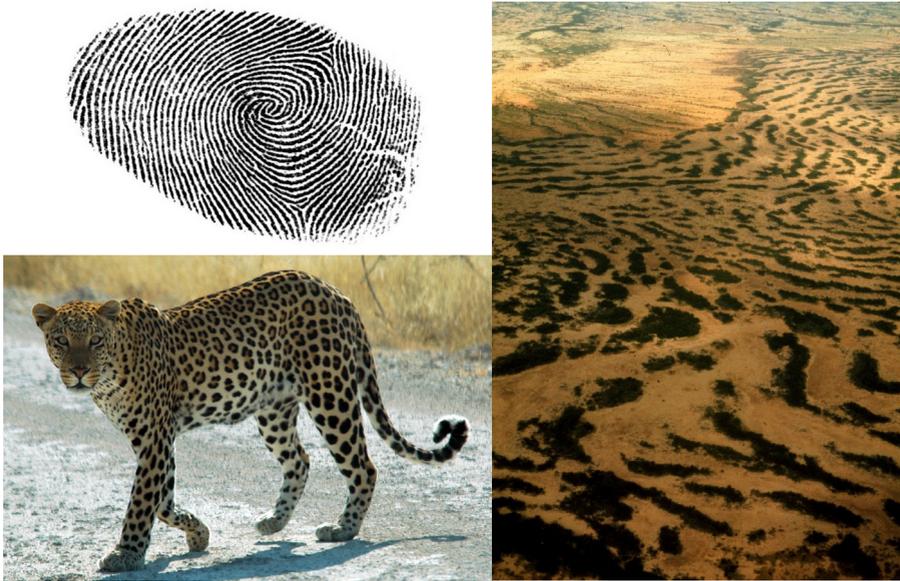


Figure 1.1: Some examples of patterns in nature: a leopard, a finger print, and vegetation patterns at the edge of the desert.

something again, etc.’, you’ll discover they are all around you. Think of a tree: its branches, twigs on the branches, leaves on those twigs are all roughly equally spaced – even the transport canals visible in the leaves exhibit a tree-like structure. On larger scales, examples of patterns are ubiquitous, even at arid places like the desert: think of wave patterns on sand dunes, or even the more or less equally spaced dunes themselves. At the edge of the desert, you can find vegetation patterns as spots and stripes. Similar spots and stripes can in turn be observed in the sky as cloud patterns.

All these patterns have an element of repetition; they can be characterised by the recurrence of a certain element. Nature is full of repeating processes: the daily cycle of the sun, the tides, the phases of the moon, the changing of the seasons. While you might be tempted to call these phenomena ‘patterns’ as well (which in a sense, of course, they are), the repetition in these phenomena is temporal rather than spatial. This distinguishes them from the patterns considered above: indeed, what mathematicians call a ‘pattern’, and we’ll stick to that from now on, is a spatial pattern. Of course, this does not mean that (temporal) change will not play a role – far from it. We’ll come to speak about ‘the dynamics of patterns’ later.

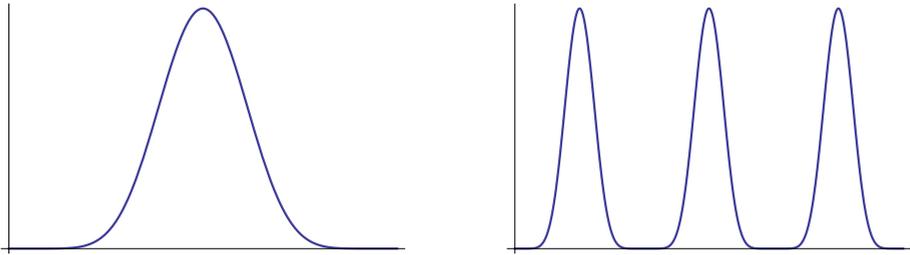


Figure 1.2: An example of a pulse and a pulse pattern.

As mentioned earlier, patterns can be characterised as the recurrence of a certain structural element, be it a spot, a stripe, a twig, a ripple, a leaf, or something else. In this thesis, the object of study is precisely such a structural element, a so-called *'pulse'*. This pulse can be viewed as a building block for more complicated patterns, see Figure 1.2. It is a natural idea to start the analysis of a pattern at its foundation, that is, by analysing its elementary building block. Once you know something (or, hopefully, a lot) about this building block, you can start to answer questions about the pattern as a whole, by looking at the ways this structural element can repeat itself. The latter, though being the obvious next step, is beyond the scope of this thesis.

Natural questions to answer when studying a pattern are: What is the repeating structural element? And how does it repeat? You can use both questions to approach the more encompassing problem of how a certain pattern is formed. The research presented in this thesis therefore falls naturally within the mathematical research area of 'pattern formation' – and in that area, within the analysis of 'localised structures'. The aforementioned pulse is an example of such a localised structure.

Depending on the pattern in question, the structural element, or localised structure, can be of more or less interest. Concerning the fingerprint, the spiral pattern itself is far more important than the narrow skin ridges of which it consists, especially in forensic analysis. In plant growth, on the other hand, the structural elements (leaves, twigs) are much more interesting. A related example is the process of embryonic limb development, which can be studied in the context of pattern formation: here, the localised structures (the arm, the fingers) are the key elements of interest. The process which causes a growing organism to develop its shape, morphogenesis, can therefore be studied in the mathematical context of pattern formation – and this is just one example of its many uses.

## 1.2.2 Dynamical systems

The mathematical techniques employed in this thesis are rooted in the field of dynamical systems, in particular that of differential equations. Without diving into the mathematics straight away, it is possible to give a flavour of how dynamical systems work, and which ideas can be used in the study of pattern formation.

A dynamical system describes the change of some quantity based on a certain system of rules. You can think of the changing position of the earth as it orbits around the sun, the concentration of chemicals when you put them together and let them react, the mass of a growing bacterial colony. There, rules governing the changes are the laws of gravity, the chemical reactions between the chemicals, and the way in which the bacterias use food and/or oxygen to reproduce. In particular, these rules can be given as a number of evolution equations. Given an initial state (a starting position, an initial concentration), an evolution equation describes how this initial state evolves in time; you might envision it like shown in Figure 1.3. Such an evolution equation is, mathematically speaking, a differential equation. An evolution equation for a quantity  $\phi$  will therefore be an equation for its time derivative  $\frac{d}{dt}\phi$ , i.e. the change of  $\phi$  in time at a certain moment:

$$\frac{d}{dt}\phi = \text{something (involving } \phi \text{ and/or } t).$$

Here,  $\phi$  plays the role of whatever your evolution equation is describing, be it temperature, an animal population, the concentration of a certain chemical, etcetera. Of course, the ‘something’ part is where the fun begins. Once you make a choice for the ‘something’-terms on the righthand side of the equation, you prescribe a certain behaviour, fixing the evolution of  $\phi$ . Different choices for the righthand side terms will lead to different dynamical behaviour of the quantity  $\phi$  – even small changes can have large consequences, as we will see in section 1.2.4.

Evolution equations are used as a model for phenomena in nature where the time evolution of certain quantities play a role, for instance the growth and decay of populations. Often, it is needed to describe in the model how the quantity in question is spread out – in space, that is. If the evolution of a quantity also depends on the way it is spread out, the space variable  $x$  plays an important role in the evolution equation describing such a process. The time evolution of a quantity  $\phi$  will depend on  $x$ , and on spatial derivatives like  $\frac{d}{dx}\phi$  and  $\frac{d^2}{dx^2}\phi$ . Therefore, such a model looks like

$$\frac{\partial}{\partial t}\phi = \text{something involving } \phi, \frac{\partial}{\partial x}\phi, \frac{\partial^2}{\partial x^2}\phi, x \text{ and/or } t.$$

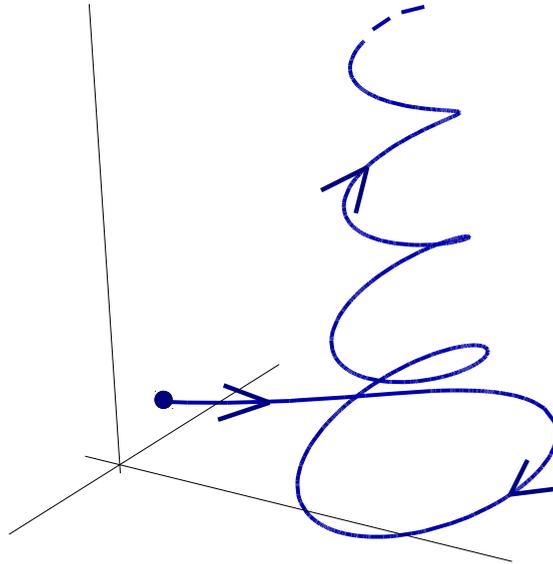


Figure 1.3: A visualisation of a dynamical system: a certain initial state evolves in time.

You may have noticed that the notation of the derivatives has slightly changed: we use ‘ $\partial$ ’ instead of ‘ $d$ ’. This notation is commonly used to emphasise that the quantity  $\phi$  depends on two variables, in this case  $x$  and  $t$ : in other words,  $\phi$  is a function of both  $x$  and  $t$ . Also, for simplicity, we stick to just one spatial variable  $x$ ; for phenomena occurring in more than one dimension, where it is for instance useful to distinguish length, width and height, more spatial variables are needed. However, the approach is often completely analogous to the one-variable case.

An important class of evolution equations where spatial spreading influences the evolution of a quantity is the class of *reaction-diffusion equations*. These reaction-diffusion equations clearly distinguish the role of the spatial derivatives of  $\phi$  ( $\frac{\partial}{\partial x}\phi$ ,  $\frac{\partial^2}{\partial x^2}\phi$ , etc.) from other terms. That is, reaction-diffusion equations are structured as

$$\frac{\partial}{\partial t}\phi = \frac{\partial^2}{\partial x^2}\phi + \text{something involving } \phi.$$

This structure can be used to explain where the name ‘reaction-diffusion’ originates. The term ‘diffusion’ means the process of spreading in space: think of a drop of milk spreading in a cup of coffee. Another everyday example is heat conduction: if you heat a pan on the stove, the heat spreads to the pan (and, more importantly, the food) due to diffusion. Mathematically, you can model diffusion most straightforwardly by a second order spatial derivative, in this case  $\frac{\partial^2}{\partial x^2} \phi$ . This term in the reaction-diffusion equation dictates how the quantity  $\phi$  spreads into space as it evolves over time.

The other terms, ‘something involving  $\phi$ ’, are called the reaction terms. The reason for using this terminology is most clear if you consider not one, but two reaction-diffusion equations, i.e. a ‘reaction-diffusion system’ (or, equivalently, you can consider a quantity  $\phi$  which has two separate components). An example is the Gierer-Meinhardt system [22], describing the evolution of the quantities  $U$  and  $V$ :

$$\begin{aligned}\frac{\partial}{\partial t} U &= \frac{\partial^2}{\partial x^2} U + V^2 - U \\ \frac{\partial}{\partial t} V &= \frac{\partial^2}{\partial x^2} V + \frac{V^2}{U} - V\end{aligned}$$

You can clearly see that the evolution of  $U$ , prescribed by the first equation, is influenced by the value of  $V$  through the term  $V^2$ . The second equation, describing the evolution of  $V$ , has in turn a term  $\frac{V^2}{U}$  which depends on  $U$ . This mutual influence can be interpreted as a reaction between  $U$  and  $V$ , which clarifies calling these terms in the evolution equation ‘reaction terms’.

Reaction-diffusion equations can therefore be characterised as evolution equations describing the spatial spreading and mutual interaction of certain quantities.

### 1.2.3 Patterns in reaction-diffusion systems

In phenomena described by reaction-diffusion models, all kinds of patterns frequently occur. This is no coincidence: the appearance of some kind of pattern in a natural phenomenon is often the incentive for researchers to try to model this phenomenon using a reaction-diffusion model. Alan Turing –the very same Alan Turing who’s famous for his groundbreaking work in computer science and for deciphering the Enigma code– was the first to postulate that (and explain why) systems of reaction-diffusion equations naturally allow the formation of patterns [50]. Patterns formed in this way are often called Turing patterns.

The most common description of the Turing patterning principle uses a so-called



Figure 1.4: Examples of diffusion: a splash of milk in coffee, a pan on the stove.

activator-inhibitor pair. Given two species (say, chemicals), we call one the ‘activator’ and the other one the ‘inhibitor’: the activator makes both species grow, while the inhibitor tries to decrease the growth of both species. This activation/inhibition description can be modelled by the reaction terms of a reaction-diffusion equation. Turing found out that if the inhibitor spreads (diffuses) much easier than the activator, a certain feedback mechanism occurs. Because of this feedback mechanism, the activator and inhibitor are not evenly spread out. Their concentration fluctuates in a very regular way, creating a pattern.

This activator-inhibitor mechanism, modelled by a system of reaction-diffusion equations, is widely believed to be the cause of a broad range of patterns in nature, such as spots and stripes on animal skins, or vegetation patterns at the edge of the desert. In Figure 1.6, you can see some examples of patterns found in a specific reaction-diffusion system (the Gray-Scott model).

What is a pattern? In the context of reaction-diffusion equations, you could say that a pattern has a clear spatial structure, so it depends on the spatial variable  $x$  in some specific way. Moreover, as in the many examples of patterns seen in section 1.2.1, you could argue that something like a pattern should not, or not really, change in time. This last condition, although it seems natural, is quite restrictive. There are clear examples of things which you undoubtedly would call a ‘pattern’, but which do move. Travelling waves (such as water waves or radio waves, or light) are examples of this: they have a periodic (spatial) structure, but move as well. Of course, you could say that when you move along with the wave, it seems to be standing still –and that’s exactly how these travelling waves are analysed in general– but that doesn’t change the fact that these waves are moving.



Figure 1.5: Alan Turing (1912-1954), pattern formation pioneer.

However, focusing on stationary patterns has some major advantages: since the pattern you're looking for is independent of time, you can imagine that its analysis in the context of reaction-diffusion equations becomes somewhat easier: there is no interplay between space and time, the pattern won't evolve. Also, this can be considered as a starting point for the analysis of patterns which do change in time. You can start to phrase questions like 'If I change some conditions, will the pattern I established start to change? Will it start to move? Will it change its shape?' In section 1.3.3, some of these questions will be addressed.

Summing up, the search for a pattern in a reaction-diffusion model may start with finding a stationary, i.e. time independent solution with a specific spatial structure to the associated system of reaction-diffusion equations. Indeed, that's exactly what part of this thesis is about: to find patterns (in particular, pulses) in reaction-diffusion systems.

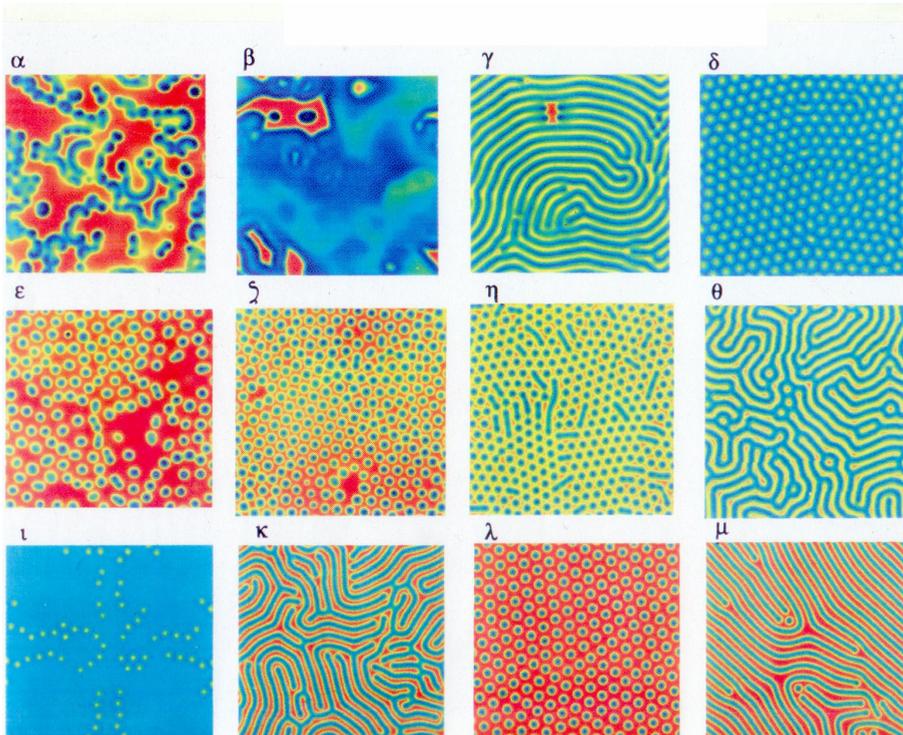


Figure 1.6: Several types of patterns in the Gray Scott model [43].

### 1.2.4 Perturbations

Let's do a little experiment. We want to study what happens to a ball when we let it drop from a certain height, say 2 metres. We can, for instance, measure how long it takes before the ball hits the ground. You can imagine that, when you perform this experiment multiple times, you won't always get the same answer. This variation in measurements can be caused by a lot of things: maybe you didn't drop the ball each time from exactly the same height, maybe you weren't always on time with your stopwatch. Those things have to do with the fallibility of the experimentalist: of course, the actual phenomenon you're investigating, the falling of a ball, isn't influenced by your incapability of measuring things exactly. Since this is a thought experiment, let's assume you are able to measure the falling time of the ball exactly.

Still, you won't get exactly the same measurement results. Maybe the ball was a little bit blown to the side by the wind, maybe the ground was somewhat uneven, maybe the atmospheric pressure changed a little, as well as the air humidity, changing the air resistance; maybe a dust speck stuck to the ball and changed its total weight, maybe the ball collided mid-air with an unsuspecting little fly, slowing it down. These are things you can't control, but which might influence the measurement results. You could argue that if you do a lot of experiments, under tightly controlled conditions, you can minimise and average out the influence of these disturbances. In the end, to explain the phenomenon of the falling ball (since we know that's just due to gravity), it shouldn't matter who's doing the experiment, or what time of day it is, or whether it's raining or not, or whether I perform the experiment in Oslo or Jakarta – but wait. That last condition actually *does* matter, although a little. We know, since the days of Newton, that the ball falls due to the mutual gravitational attraction between the ball and the Earth. When you start to measure the gravitational acceleration at different places on Earth (for instance, by performing our falling ball experiment), you'll find out that this gravitational acceleration  $g$  differs a little bit from place to place. Approximately,  $g = 9.8 \text{ m/s}^2$ ; in Oslo, we have  $g_{\text{Oslo}} = 9.825 \text{ m/s}^2$ , while in Jakarta, we have  $g_{\text{Jakarta}} = 9.777 \text{ m/s}^2$ . While the other small influences on the falling ball were random, and didn't have anything to do with the physics underlying the phenomenon of the falling ball, the location on Earth introduces a systematic, though small, change.

When a researcher studies a natural phenomenon, he or she tries to establish which processes or laws are really instrumental to understanding the phenomenon, and which processes are just noise, disturbances. If you would write down a model, i.e. an equation for the falling ball, you wouldn't incorporate the time of day or the color of your eyes, since you know those don't matter. Also, you wouldn't incorporate the wind or the air humidity, since you know those aspects of nature are not underlying the phenomenon of the falling ball. Although they might influence its movement a little, that's not what you want to study in the end. In other words, you want your equations (your model, your physical laws) to be as clean, as simple, as possible. This is one of the reasons why it's hard to write down a good model as a scientist: you need a lot of knowledge and experience to single out the processes that really matter for the phenomenon you're studying.

As for the falling ball, we can write down a formula for its falling time using Newton's second law, incorporating only gravity. If we call the falling time  $t$  and the gravitational acceleration  $g$ , this gives (for a falling height of 2 metres, and forgetting

about units)

$$t = \frac{2}{\sqrt{g}}$$

Since the value of gravitational acceleration  $g$  differs from place to place, the falling time differs from place to place. Of course, if we would just want an approximate result for the falling time, we could always use the approximate value for  $g$  – even though it’s not exact, the ‘real’ value of  $g$  (and therefore the real value of  $t$ ) is not far from it. Using the average value  $g_{\text{average}} = 9.81 \text{ m/s}^2$ , we obtain  $t_{\text{average}} = 0.64 \text{ s}$ . Now we can compare the actual values of the gravitational constant at Oslo and Jakarta with the average value:  $g_{\text{Oslo}} = g_{\text{average}} + 0.015 \text{ m/s}^2$  and  $g_{\text{Jakarta}} = g_{\text{average}} - 0.033 \text{ m/s}^2$ . In this way, we can write the gravitational acceleration anywhere on Earth as  $g = g_{\text{average}} + \varepsilon$ , where the value of  $\varepsilon$  depends on where you are. Also, as we’ve already seen,  $\varepsilon$  is quite small compared to  $g_{\text{average}}$ . If we incorporate this in our equation for the falling time, we get

$$t = \frac{2}{\sqrt{g_{\text{average}} + \varepsilon}}$$

This formula, giving the falling time of a ball dropped from a height of 2 metres anywhere on Earth, is an example of a model with a *perturbation*. In this way, you can immediately read off a number of aspects characteristic to the falling ball phenomenon. For instance, if you ignore the small variations in  $g$ , setting  $\varepsilon = 0$ , you can immediately see how to calculate the approximate, average falling time  $t_{\text{average}}$ , namely as

$$t_{\text{average}} = \frac{2}{\sqrt{g_{\text{average}}}}$$

Moreover, you can see that, as long as the perturbation  $\varepsilon$  is small, the falling time won’t differ very much from the average falling time, see Figure 1.7. This last characteristic, that small changes in the model have small effects on its outcomes, is the defining property of so-called *regular* perturbations.

The opposite is true for *singular perturbations*, where small perturbations in the model can have large effects on the quantities described by the model. This sounds somewhat counterintuitive, but there are everyday examples where singular perturbations play an important role.

Singular perturbations are almost always associated with sudden changes, or fast transitions. A good example is heat conduction, which we’ve encountered before in section 1.2.2, where it was intimately connected to the term ‘diffusion’.

If you put a pan on the stove, the heat from the stove spreads very quickly through the

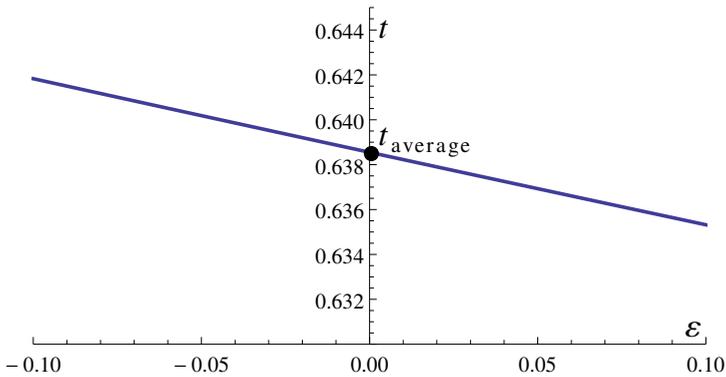


Figure 1.7: The value of the falling time  $t$  for small perturbations  $\varepsilon$ . As you can see from the axis scaling, the falling time doesn't change that much when  $\varepsilon$  is small.

metal of the pan: this metal conducts heat very well. If the pan would be ceramic, this would be completely different: since ceramic materials are quite good at insulating heat, the pan would heat up very slowly, since the heat from the stove would hardly spread through the ceramic material. The difference between metal and a ceramic material such as porcelain in terms of heat conduction can be seen clearly by comparing Figures 1.8 and 1.9. In the ceramic material, the heated spot in the middle does not spread out like in the metal. Therefore, there is a sharp transition between the heated region and its surroundings: at the edge of the heated region, there is a sudden drop in temperature. That this phenomenon is intimately related to singular perturbations, becomes clear once you consider the model underlying the phenomenon of heat flow through materials.

Heat spreads through materials due to diffusion. Indeed, the spread of heat through a material can be modelled by a very simple evolution equation known as the 'heat equation':

$$\frac{\partial}{\partial t}\phi = \alpha \frac{\partial^2}{\partial x^2}\phi$$

where, in this case,  $\phi$  is the temperature at a given place in the material at a certain time. This is a very basic reaction-diffusion equation, or rather just a diffusion equation since there are no reaction terms present (see section 1.2.2). Here,  $\alpha$  is called the thermal diffusivity: it is a constant, and its value depends on the material considered. A material which doesn't transfer heat very well has a very small thermal diffusivity. To make a connection with the previous example, let's call this small thermal diffu-

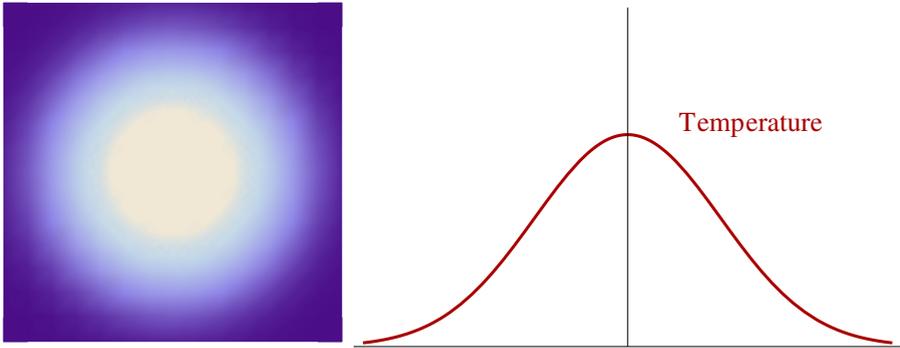


Figure 1.8: A heated spot on a metal pan: you can clearly see that the heat diffuses very well.

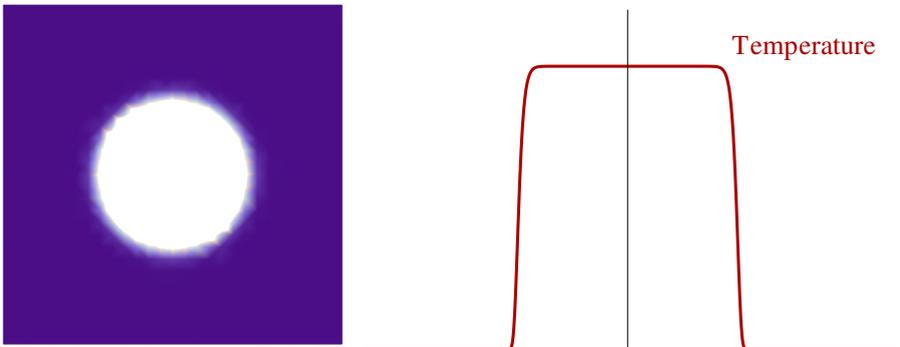


Figure 1.9: A heated spot on a porcelain pan: the heat hardly diffuses, and there is a sharp temperature transition from the heated spot to the outside.

sivity constant ‘ $\varepsilon$ ’, such that the heat equation for such a highly insulating material becomes

$$\frac{\partial}{\partial t}\phi = \varepsilon \frac{\partial^2}{\partial x^2}\phi$$

Like in the falling ball example, you can ask yourself what happens when we neglect the small  $\varepsilon$  term, i.e. set  $\varepsilon = 0$ . In this case, something very drastic happens: the heat equation simplifies to

$$\frac{\partial}{\partial t}\phi = 0$$

In other words, the temperature  $\phi$  does not change. This means that the transition

between the heated spot and the surrounding areas really is a sharp transition: the heated spot stays heated, since it's insulated by the surrounding material, and the surrounding material stays cold. Of course, this doesn't give a completely realistic description of the situation: in reality, the heat will slowly spread and the heated spot will cool down slowly. However, this so-called 'singular limit' gives a quite convenient approximation of the real phenomenon, as long as you're willing to overlook the fact that a perfectly insulating material can hardly be expected to exist.

The limit  $\varepsilon = 0$  is called 'singular', because it throws away a term in the model which is crucial to the description of the phenomenon, in this case heat diffusion. This tendency of singular limits to dispose of instrumental terms is a recurring theme in the study of singular perturbations, and can often be used to the advantage of the researcher, since it leads to a vast simplification of a possibly complicated model. You can think of it as a trade-off: by setting  $\varepsilon = 0$  in a singular limit, it suddenly becomes possible to solve some equations, since the complicated system is severely reduced. On the other hand, you've thrown away quite a lot: it is often unclear how the analysis in the singular limit can tell you something about the case when  $\varepsilon$  is not equal to zero (but still very small). In the above example of heat diffusion, we already understood the phenomenon which was being modelled by the 'full' heat equation, so we could interpret the singular limit. In other applications, this is often not so easy.

We've seen that, in the case of very slow heat diffusion, there is a sharp transition between the heated spot and the surrounding areas. In reality, this transition is not discrete as in the singular limit, but occurs very fast. To obtain a better understanding of what's happening at the transition, it's a good idea to zoom in on that transition zone. If you do that, you'll see a gradual change from high to low temperature, but now gradual on a very small spatial scale. The best way to describe the temperature distribution in the porcelain pan is as follows:

1. Start far away from the heated spot. There, the temperature is low. If you start to 'walk' towards the heated spot, nothing much happens: the temperature stays the same.
2. Suddenly, there is a huge jump in temperature. You're now at the edge of the heated spot. To obtain a better view of what's happening here, you will have to slow down and zoom in at the transition zone: you'll see that, on this small scale, the temperature steadily rises.
3. After the jump, you're in the heated zone. Again, nothing much happens: the high temperature is everywhere more or less the same.

This phenomenon of scale separation, where the best way to describe a situation is to analyse different parts of it at different scales, occurs naturally when singular perturbations are present. The general approach in such situations is therefore to separate the problem into different scales, analyse each scale separately, and then try to paste these descriptions together to obtain a consistent global picture. This idea of separating scales, analysing reduced problems and trying to combine the results will be explained in more detail in section 1.3.1.

### 1.2.5 What is this thesis about?

We've come to a point in this introduction where it's possible to understand the title of this thesis: 'Pulses in singularly perturbed reaction-diffusion systems'. This thesis is about the analysis of a certain pattern, namely a single pulse, in the context a certain class of models, namely reaction-diffusion systems. Moreover, these reaction-diffusion systems possess a certain very useful quality: they are singularly perturbed. In section 1.2.1, we've seen what patterns are and how mathematicians think of them; in section 1.2.2, it was explained what we mean by 'reaction-diffusion systems', and in section 1.2.3, it was indicated why patterns can occur in these reaction-diffusion systems, and how they are characterised mathematically. The concept of 'perturbations', and in particular 'singular perturbations', has been introduced in section 1.2.4.

The concepts introduced in the previous sections will be used to clarify the research methods presented in the upcoming sections. I'll give an overview of the methods used in my research, and present the general approach to analysing patterns in reaction-diffusion systems.

## 1.3 Methods

When you start to analyse pulses (or any pattern, for that matter) in reaction-diffusion systems, you start by asking the question 'Does a pulse exist at all?'. In other words, you start to investigate whether the reaction-diffusion system you're considering admits something like a pulse solution. This is the question of *existence*, which will be addressed in section 1.3.1. If the system doesn't admit a pulse solution, there's not much to investigate. On the other hand, if the system does admit the pulse solution you're looking for, you're not done yet. The next obvious question is (and I'll explain why this question is obvious) 'Is this pulse stable?'. The notion of *stability* has not been introduced, but will be explained in detail in section 1.3.2. The short answer to

the question why stability is important, is that if a pulse exists but is unstable, you often won't observe it in 'real life' applications. Therefore, a priori, the stable pulses are the ones you're looking for.

However, that's not the whole story. Disregarding unstable pulse solutions as wholly uninteresting doesn't do them justice in all cases. There are numerous examples of exciting phenomena which can unfold when you look at an unstable pulse. Therefore, you should ask the question 'If the pulse is unstable, what will happen?'. This will open the gateway to a vast realm of possibilities in the field of *dynamics*, i.e. the motion of a pulse. I'll explain how you could go about analysing one of those many possibilities in section 1.3.3. This is also the subject of the last chapter of this thesis, chapter 4.

In Figure 1.10, you can find an overview of this research approach.

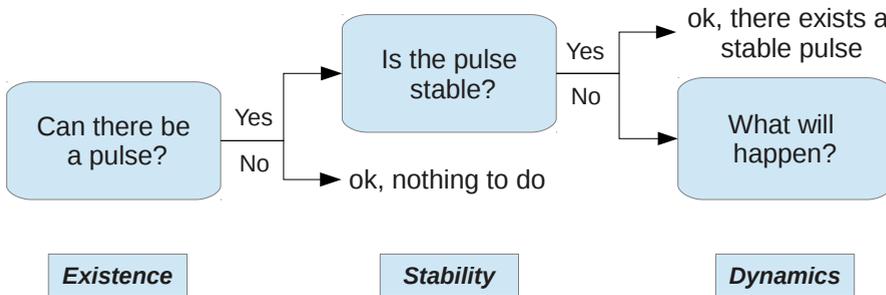


Figure 1.10: General research approach.

### Parameters

In most cases, the answer to both the existence and stability question is a lot more subtle than just yes or no. This is the case when the reaction-diffusion system you're looking at depends on *parameters*. Parameters are constants, which in other words don't change in time or in space, which you can therefore choose freely. Parameters often tell you something about the environment in which the phenomenon you're modelling takes place, about the conditions you're dealing with. A good example is the heat equation from section 1.2.4,

$$\frac{\partial}{\partial t}\phi = \alpha \frac{\partial^2}{\partial x^2}\phi$$

The thermal diffusivity constant  $\alpha$ , which depends on the material you're looking at, is an example of a parameter. The advantage of using parameters is that you don't have to redo your entire analysis if you decide to switch from analysing one material to analysing another. Since the parameter  $\alpha$  is fixed from the model point of view (although its precise value may not be specified), you can do your analysis for all values of  $\alpha$  in one go. You'll obtain an answer (for instance, to the question whether a pulse exists) which still depends on  $\alpha$ , of course; then, you can see how that answer changes when you pick different values of the parameter  $\alpha$ .

In general, a reaction-diffusion system always has some parameters, whose value you can pick at your own leisure, depending on the context you're working in. Therefore, the existence question then is not really *if* but *when* there exists a pulse, or 'For which parameter values does a pulse exist?'. Exactly the same situation occurs for the stability question. Say you've chosen suitable parameter values such that a pulse exists (you can still have a large number of ways to do that), then the question arises *when* your pulse is stable, i.e. for which parameter values. Within the set of parameter values for which your pulse exists, you can make a division between parameter values for which the pulse is stable, and for which it is unstable; see Figure 1.11 for a visualisation.

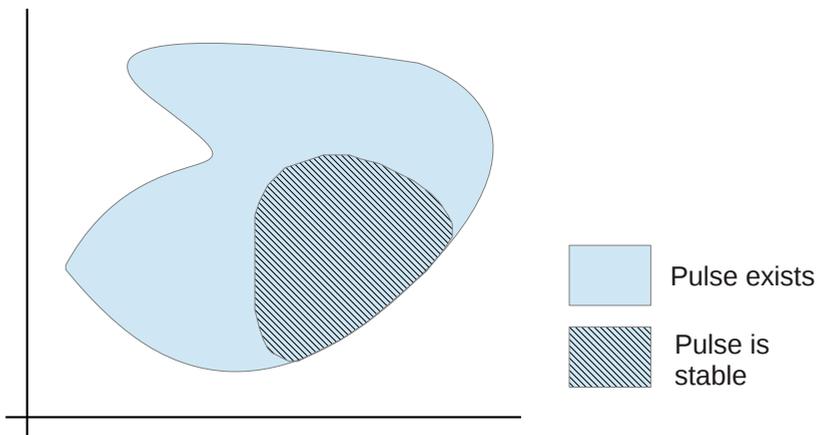


Figure 1.11: The existence and stability regions in parameter space.

## The parameter $\varepsilon$

In this thesis, you'll come across a large number of parameters. There's one parameter that is the most important of them all, and that's  $\varepsilon$ . Everywhere in this thesis,  $\varepsilon$  is a very small parameter. It acts as a perturbation, just as in section 1.2.4, and is instrumental in obtaining virtually all results in this thesis. Without the help of the small parameter  $\varepsilon$ , the research presented in this thesis couldn't have been carried out. The reason why  $\varepsilon$  is so helpful in the analysis of the existence and stability of pulses was already revealed a little bit in the context of singular perturbations. In the following sections, the fact that  $\varepsilon$  is very small is used extensively.

Of course, how small 'small' actually is, depends on the situation. In the example of the falling ball, the deviation of the gravitational acceleration in Oslo and Jakarta from its global average value was very small – that is, very small when compared to that average value: the deviations were something like  $\pm 0.03 \text{ m/s}^2$ , while the global average gravitational acceleration was  $g_{\text{average}} = 9.81 \text{ m/s}^2$ .

Since we're dealing in this thesis with quite general reaction-diffusion equations and, more importantly, with parameters whose value is not specified, the question how small  $\varepsilon$  needs to be to obtain sensible results is a hard one. However, once you've chosen your parameter values, you can often answer this question. Therefore, the most important results in this thesis, which are presented as Theorems, start with the phrase 'Let  $\varepsilon$  be small enough'.

The general idea is that, when  $\varepsilon$  is small enough, you can use the information obtained for the case  $\varepsilon = 0$  to prove results about the case when  $\varepsilon$  is not zero, but small. This often works very well because the analysis for case  $\varepsilon = 0$  is most of the time much simpler than for nonzero  $\varepsilon$ . Again, you can prove several results when  $\varepsilon$  is small enough. Although this can sound restrictive, experience tells us that in practice, these results continue to hold for surprisingly large values of  $\varepsilon$ .

In the following sections, you'll see how  $\varepsilon$  is being used to obtain results about the existence, stability and dynamics of pulses in singularly perturbed reaction-diffusion equations.

### 1.3.1 Existence and construction

As we already noted, the first question you ask yourself when analysing pulse solutions is 'Do these pulses exist at all?'. Questions about existence of certain things are not uncommon in mathematics. However, when you're working in a very abstract context, it's often not possible to go very much beyond that. You might be able to prove that something (in our case, a pulse) has to exist, but where you can find it and

what it looks like is often not known, or even not knowable.

In the case of pulses in reaction-diffusion systems, that's very different. The way the existence of pulses in these models is proved is *constructive*. That means that we start looking for a pulse, and once we find it, we know what it looks like. In other words, we obtain an explicit expression which describes our pulse in an approximate way, and that's often more than you could hope for. If you've got an explicit expression, you can make plots, investigate the specific shape of the pulse and draw several conclusions. That's the clear advantage of a so-called constructive proof of existence. You not only know if (and when) a pulse exists, you have it at your fingertips at the same time.

### Pulse shape

What is a pulse? The shortest answer is that it's a function whose graph looks like the one in Figure 1.2. It has a single hump, and gradually decreases to the left and to the right to a constant value, most often to zero. That last property shall be the key to identify pulses. You can see that, if you're far away from the hump, that the graph of the function is very flat and very close to zero. In other words, when you get further and further away from the hump, the function and its slope (its derivative) should get closer and closer to zero.

Let's draw the function in another way. In Figure 1.12, the function is graphed in such a way that you can see its function value and the value of its slope. The horizontal axis gives the function value, the vertical axis the slope of the function. In this picture, you can't really see the spatial variable  $x$  anymore. Therefore, at some points, I indicated what the corresponding  $x$ -value is.

There are some things to notice in this picture. First, the pulse is now a kind of loop which starts and ends where the axes meet, at the origin. Here, both the function value and the value of the derivative are zero. Also, you can see that the loop doesn't 'really' close: as  $x$  becomes larger and larger (or more and more negative), the function value and its derivative come closer and closer to zero, but they never really reach it. However, you can get as close to the origin as you want, by going to large enough values of  $x$ . Moreover, you can argue why a pulse should be equivalent to such a loop: when you start far to the left, the function value will increase, so its slope will be positive. As you come close to the top of the hump, the pulse becomes flatter and flatter, until you're at the top, where the derivative is zero. Then you start going down, i.e. the slope is negative, while the function value decreases towards zero. The function becomes increasingly flat, so the derivative, while still negative, also becomes very small.

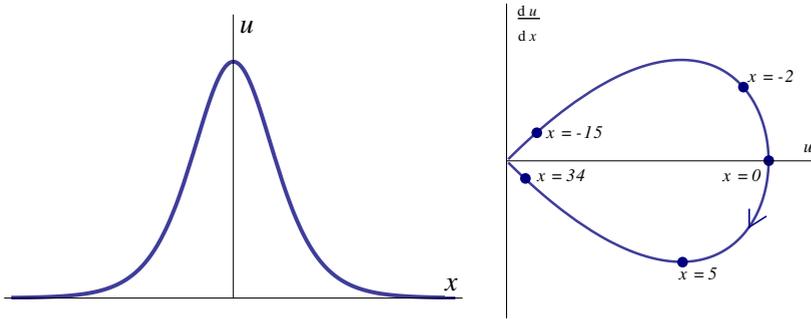


Figure 1.12: An example of a pulse: left, as a function of  $x$ ; right, as a loop in the phase plane. The pulse height is given by  $u$ , its slope by  $\frac{du}{dx}$ .

A plot such as shown in Figure 1.12 is called a phase plane plot, and is often used to clarify the behaviour of two-dimensional dynamical systems. The loop is an example of a so-called ‘orbit’ in such a dynamical system. If you pick a point on this loop as an initial value and let the dynamical system run, you’ll start to move along this loop, thereby tracing out an orbit. This also holds when you move backwards, then you would trace the loop the other way around. In Figure 1.12, the direction you’ll move in if you let  $x$  increase, is indicated by an arrow.

### Four-dimensional problem

Back to our original problem, the construction of a pulse in a reaction-diffusion system. In this thesis, I look at reaction-diffusion systems which have two components, so which can be written as two reaction-diffusion equations. The components will influence each other through the reaction terms. An example of such a system is the Gierer-Meinhardt model, as mentioned in section 1.2.2. I’ll also call these components  $U$  and  $V$ . Also, I’ll introduce a small parameter  $\varepsilon$  in the evolution equation for  $V$ , which will come in handy at a later stage. A general reaction-diffusion system having these properties looks like

$$\begin{aligned} \frac{\partial}{\partial t} U &= \frac{\partial^2}{\partial x^2} U + F(U, V) \\ \frac{\partial}{\partial t} V &= \varepsilon^2 \frac{\partial^2}{\partial x^2} V + G(U, V) \end{aligned}$$

It doesn’t really matter what the reaction terms  $F(U, V)$  and  $G(U, V)$  are; all that’s important, is that they make sure that  $U$  influences the evolution of  $V$  and  $V$  influences

the evolution of  $U$ . In other words, these equations are coupled. You cannot solve one without the other, you'll have to analyse them simultaneously.

If we start looking for a pulse in this system, we're looking for a stationary solution of a particular form, i.e. something which does not change in time. That means that our equations simplify somewhat:

$$\begin{aligned} 0 &= \frac{d^2}{dx^2}u + F(u, v) \\ 0 &= \varepsilon^2 \frac{d^2}{dx^2}v + G(u, v) \end{aligned}$$

You'll notice that I use lower case letters  $u$  and  $v$  instead of upper case ones: this is to emphasise that we're looking for something which does not depend on time, only on the spatial variable  $x$ . In other words, both  $u$  and  $v$  are functions of just  $x$ . For that reason, the notation of the derivative also changed a little: instead of ' $\partial$ ', there's now a ' $d$ '. This indicates that the only variable we're need to worry about is  $x$ .

If we want to obtain a pulse solution for this system, we want *both* components  $u$  and  $v$  to look like Figure 1.2. If we draw them in one picture, this would look like Figure 1.13. Here, the  $u$ -component is indicated in blue and the  $v$ -component in red. We can try and draw this in terms of the 'loop' picture introduced earlier. However, that's a little problematic: since we've got two components, and for each component we need two axes (one for the component itself, one for its derivative), we'll need  $2 \times 2 = 4$  axes in total. That means we'll have to work in a four-dimensional space. Mathematically, that's absolutely no problem: you just start to work with four different components. Visualisation-wise, this introduces large difficulties: how can you picture something in four dimensions? Well, in this case, I don't think you really can. However, I'll try to convey some ideas using three-dimensional plots, which should help you get the complete picture of what's happening.

In any case, we want to do something similar to the phase plane plot as shown in Figure 1.12. Our pulse in both components is now such a loop, but hanging in four-dimensional space instead of lying in the two-dimensional phase plane. Because the equations are still coupled to each other, we cannot analyse the equations for  $u$  and  $v$  separately: it's no use to draw to separate phase planes (one for  $u$ , one for  $v$ ) and combine the results – at least, that's what you would think! It will turn out that, in our case, it *is* possible to separate the  $u$ - and  $v$ -components. The reasons why this wholly unexpected possibility arises will become clear in a moment.

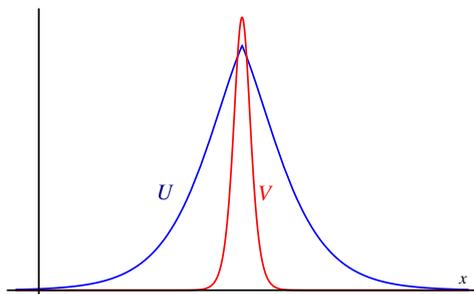


Figure 1.13: A pulse in both the  $u$ - and the  $v$ -component.

We're looking for a pulse in two components,  $u$  and  $v$ ; for both components, it is true that their function values and their slope come close to zero when you're far away from their humps. That means that our loop hanging in four-dimensional space is connected with its tip to the place where all the four axes meet, i.e. where  $u$ , its derivative,  $v$  and its derivative are zero. Figure 1.14 gives an idea of what's happening. I omitted the axis for the derivative of  $v$  – when plotting a four-dimensional picture in three dimensions, you have to make some choices. Remember: we're still not sure if and when such a loop exists, that's exactly what we're trying to find out.

What are the defining characteristics of this loop? It starts at the origin of the axes, makes some excursion through the four-dimensional space, and then returns to the origin. This view is going to help us establishing the existence of such a loop. We're going to adapt a dynamical-systems point of view, and start looking for an orbit which:

- a) goes towards the origin as  $x$  becomes very large, and
- b) goes towards the origin as  $x$  becomes very negative.

We proceed as follows: let's look at *all* the orbits which a) go towards the origin as  $x$  becomes very large, and bundle them together. This bundle is an example of a manifold. For the purpose of this introduction, you can think of something resembling a sheet of paper. We do the same with all the orbits which b) go towards the origin as  $x$  become very negative; in this way, we obtain another bundle. Now, because the loop orbit we're looking for goes towards the origin both as  $x$  becomes very large and as  $x$  becomes very negative, we see that this orbit must belong to both bundles. If you

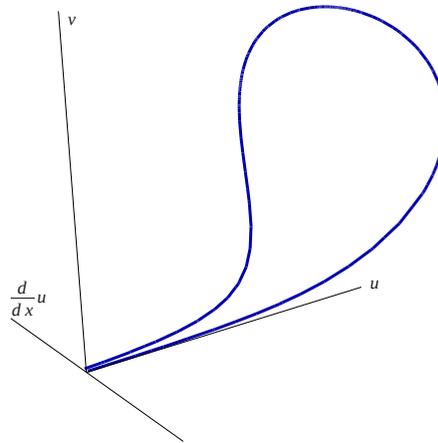


Figure 1.14: An impression of a loop hanging in four-dimensional space. The axis for the derivative of  $v$  is omitted.

picture these bundles as sheets, that means that such a loop must be in both sheets. This is only possible if these sheets intersect in some way. This situation is depicted in Figure 1.15.

We've now reformulated the question 'Does there exist a pulse solution?' as 'Do these two sheets intersect?'. Indeed, these two questions are equivalent: if there is a pulse, then the two sheets should intersect. Vice versa, if the two sheets intersect, then there must be a pulse. The problem of the existence of the pulse is now phrased in geometric terms. Apart from visualisation purposes, this approach has its advantages in other aspects of the problem. Combined with the fact that our problem is singularly perturbed, this will lead to a complete understanding of the pulse existence problem, and ultimately to its solution.

### Scale separation

We now turn to the important property our reaction-diffusion system has: it is singularly perturbed. A small parameter  $\varepsilon$  can be found in the evolution equation for  $V$ :

$$\frac{\partial}{\partial t} V = \varepsilon^2 \frac{\partial^2}{\partial x^2} V + G(U, V)$$

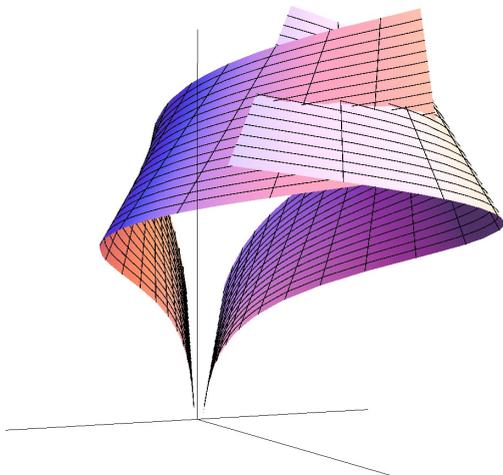


Figure 1.15: Intersecting sheets.

Just as in the example of the heat equation, this means that there will be a steep transition, in this case in the  $V$ -component. For the pulse we're looking for, this means that the pulse in the  $v$ -component will become very narrow and very steep, as you can see in Figure 1.16. In other words, the  $v$ -component of the pulse is practically everywhere just flat, almost equal to zero, except in a very small zone, where it's sharply peaked. That means that, apart from that very thin peak, we can just treat  $v$  as if it were zero. That means that everywhere except for that very small zone, the  $u$ -component can be described with

$$0 = \frac{d^2}{dx^2}u + F(u, 0)$$

This is much easier than before: the  $u$ -equation does not depend on  $v$  anymore, because  $v$  vanishes almost everywhere. In other words, the system decouples. We can analyse this system using the phase plane, where the horizontal axis gives the value of  $u$  and the vertical axis gives the value of its derivative  $\frac{d}{dx}u$ . Remember, we want to have a pulse in both components, so in particular in the  $u$ -component: that means that we're looking for orbits that go towards the origin as  $x$  becomes very large or very negative. Since  $v$  is already practically flat and close to zero except for a very small zone, we just need to focus on  $u$  and find orbits in the phase plane belonging to the

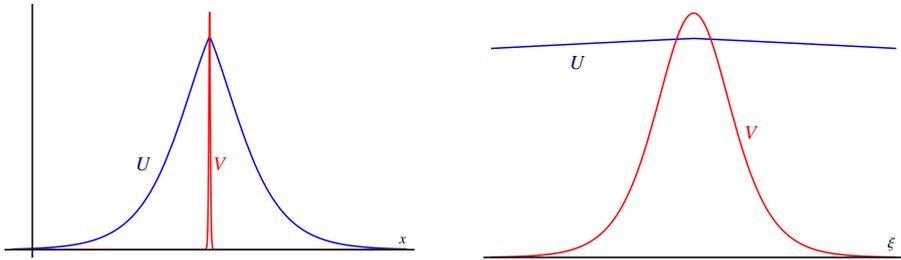


Figure 1.16: The  $u$ - and  $v$ -components of the pulse when  $\varepsilon$  is very small. The pulse in the  $v$ -component is very narrow and very steep. To the right, a zoom of the left picture near the narrow and steep  $v$ -pulse. Here, the  $u$ -component hardly changes at all.

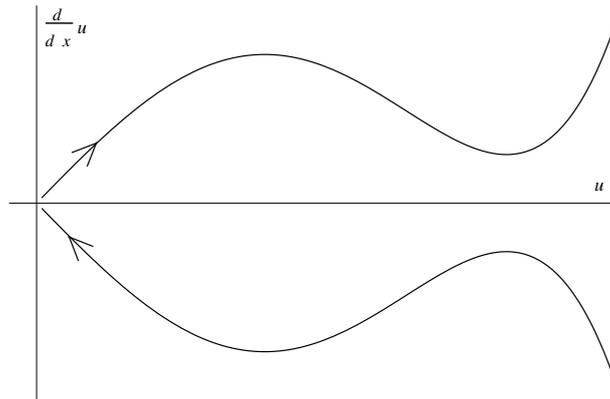


Figure 1.17: A possible phase plane, with orbits going towards and moving away from the origin.

$u$ -equation which go towards the origin in that phase plane. A possible phase plane situation where this happens is sketched in Figure 1.17.

This reduction, where we can treat  $v$  as if it were zero, holds almost everywhere, except in a very small zone where the narrow  $v$ -peak is. Just as in the example of the heat equation, we'll try to zoom in on that small zone, to find out what's happening. That means we will try to use a very small spatial scale to rewrite our system. If

we were measuring space in metres, now we're going to millimetre scale, or even smaller. We can do that by rescaling our spatial variable  $x$ . We'll use our favourite small parameter  $\varepsilon$  for this, and introduce the new spatial variable  $\xi$ . This new spatial variable is just a rescaling of the old variable  $x$ , namely  $\varepsilon\xi = x$ . If  $\varepsilon = 0.001$ , this would indeed mean that if  $x$  is measured in metres, then  $\xi$  would be measured in millimetres.

Using this small scale variable to zoom in to the  $v$ -peak, we see that on this small spatial scale, the  $u$ -component hardly changes at all, see Figure 1.16. Just as we treated  $v$  as if it were zero outside this small zone, we can treat  $u$  as if it were constant inside this small zone. Then, the equation for  $v$  becomes

$$0 = \frac{d^2}{d\xi^2}v + G(u = \text{some constant}, v)$$

You'll immediately notice that the small parameter  $\varepsilon$  is gone, this is incorporated in the fact that we're looking at a small  $\xi$ -scale now. Also, the  $v$ -equation does not really depend on  $u$  anymore – well, of course it still does, but not in a very complicated way, since we decided to treat  $u$  as if it were constant on this very small interval. In other words, we can treat the value of  $u$  as a *parameter* on this small interval. For whatever constant value of  $u$  is appropriate (we'll see how to choose that appropriate value later), we can now perform our phase plane trick on this  $v$ -equation. In our small spatial scale  $\xi$ , the sharp narrow  $v$ -peak looks like a 'regular' pulse, just like the one in Figure 1.2. Therefore, its phase plane, drawn in Figure 1.18, looks very much like that in Figure 1.12.

### Geometric theory

It is clear that the small parameter  $\varepsilon$  has been very instrumental in pulling apart the  $u$ -equation and the  $v$ -equation, allowing us to analyse them separately. Now the question is: how do we combine these two pictures? How do we use our insight on the reduced system outside the very small area, where only  $u$  is interesting, and the other reduced system inside the very small area, where only  $v$  is interesting, to build up a complete picture for the complete system?

Here, the previously developed picture of the intersecting sheets returns. Without going into details, it turns out that you can use your knowledge about the reduced systems on the two different spatial scales to describe what these sheets look like. Since the extreme situation, where  $\varepsilon = 0$ , can be very well understood using both reduced equations (one for  $u$  on the 'normal'  $x$ -scale, one for  $v$  on the small  $\xi$ -scale), the sheets are also very well understood in this singular limit. Now it's time to in-

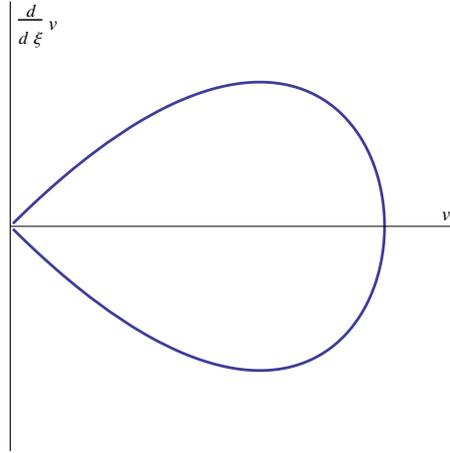


Figure 1.18: The  $v$ -peak, as depicted in the phase plane.

voke something called ‘geometric singular perturbation theory’, which was devised by Neil Fenichel in the 1970s [18, 19], to conclude something about the situation where  $\varepsilon$  is not zero, but very small.

Fenichel tells us that geometric objects like those sheets we introduced, change very little in shape when you go from the situation when  $\varepsilon = 0$  to the situation where  $\varepsilon$  is not zero but small. In other words, when you ‘turn on’ the small parameter  $\varepsilon$ , the sheets will be a little bit perturbed, but not much. That means that if the sheets intersected properly (in technical terms: ‘transversally’) when  $\varepsilon = 0$ , they still intersect when  $\varepsilon$  is small but not zero. In other words, you can use the reduced equations for  $u$  and  $v$  to obtain a picture like Figure 1.15, which still holds when  $\varepsilon$  is small but not zero. You can imagine why  $\varepsilon$  must be small: if it becomes too large, the sheets will deform too much and therefore maybe not intersect anymore, see Figure 1.19.

By combining the information from the phase plane of the  $u$ -equation with information from the phase plane of the  $v$ -equation, it is possible to determine when the two sheets intersect. Since in the  $v$ -equation, the value of  $u$  was treated as a constant and therefore acted as a parameter, this ‘intersection criterion’ determines which values this  $u$ -value can take such that the sheets intersect. This determines the height of the tip of the  $u$ -pulse. Once you’ve chosen this (almost) constant value of  $u$ , the phase plane of the  $v$ -equation (see Figure 1.18) is fixed, and therefore the loop orbit in that phase plane (which represents the  $v$ -pulse) is also fixed.

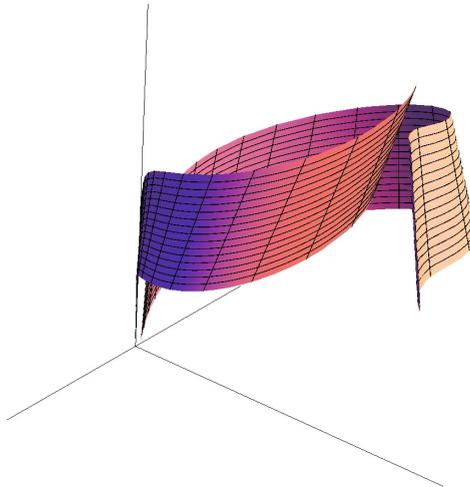


Figure 1.19: Non-intersecting sheets.

This way, a geometric approach can be used to construct a pulse in a singularly perturbed reaction-diffusion system, thereby establishing its existence. Moreover, we know approximately what it looks like: the  $u$ -pulse looks very much like the orbits in the phase plane for the  $u$ -equation, Figure 1.17, which go towards the origin as  $x$  becomes very large or very negative. In turn, the  $v$ -pulse looks very much like the loop orbit in the phase plane for the  $v$ -equation, Figure 1.18.

### 1.3.2 Stability

To introduce the concept of stability, let's return to our example of the falling ball. When you start to think about the experiment in practical terms, you'll find out that there are a large number of unwelcome circumstances which might, and often will influence your measurement of the falling time. Environmental disturbances like wind, changing air humidity and atmospheric pressure, and the occasional intervening fly will certainly influence the path of the ball by nudging it a little bit off course, slowing it down or speeding it up. Even though you can try to eliminate these disturbances by executing your experiment under tightly controlled conditions, you know these conditions will play a role once you start to consider falling balls or falling objects in nature.

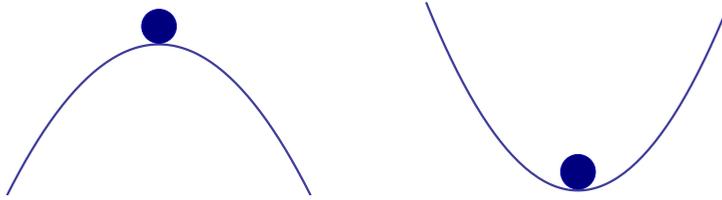


Figure 1.20: Stability visualised.

Evolution equations like reaction-diffusion systems are widely used in natural sciences like biology, chemistry, physics, geology and ecology to model natural phenomena. These models try to capture the ‘essence’ of the driving forces behind these phenomena, disregarding environmental noise. However, in the real world, such noise will always be present. Therefore, if you’re looking for a pattern in your model of which you hope you’ll observe that same pattern in nature, it’s important to know whether this pattern is robust under such naturally occurring disturbances. If even the smallest nudge will cause the pattern to change into something else, you’ll never have the chance to observe that pattern in nature.

The concept of stability is maybe best illustrated by a ball on top of a hill: if you give the ball a small nudge, it will start rolling down the hill. Therefore, the situation where the ball rests on top of the hill is unstable. On the other hand, if your ball lies at the bottom of a pit, giving the ball a small nudge won’t change much: the ball will roll around a little bit, and after a while come to rest at the bottom of the pit again. Therefore, the situation where the ball rests at the bottom of a pit is stable. For an illustration, see Figure 1.20.

Both situations describe a certain equilibrium. The equilibrium of the ball on top of the hill is an unstable one, while the equilibrium of the ball at the bottom of the pit is stable. You can determine the stability of these equilibria by considering a situation where the ball has been given a small displacement. If that small displacement starts to grow (in the case of the ball rolling down the hill), the equilibrium is unstable; if the small displacement diminishes and eventually disappears, the equilibrium is stable. When you want to determine the stability of an equilibrium, you therefore always start by considering a situation ‘nearby’, and see what will happen.

How can we translate this to the stability of patterns, or in particular, the stability of our pulse? How can we take something to be ‘nearby’ our pulse, and how will we discover what will happen for such a ‘nearby’ configuration? To give you an idea how to approach this, I’ll first have to say something about linearity.

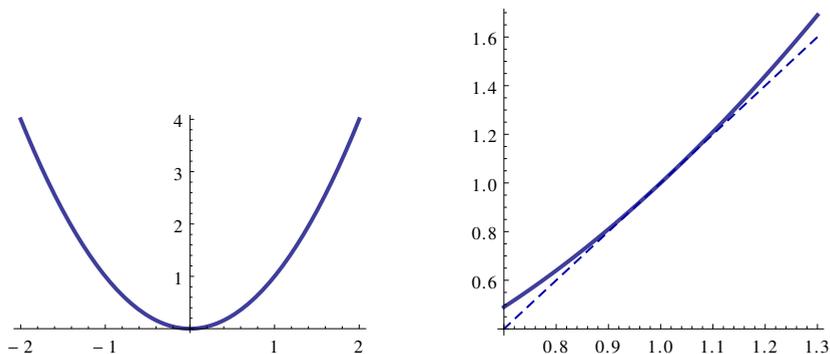


Figure 1.21: The graph of  $x^2$ , zoomed in near  $x = 1$ . As you can see, the dashed tangent line at  $x = 1$  approximates the function very well near  $x = 1$ .

### Linearity

Let's consider a simple function, like  $x^2$ . As you probably know, the graph of this function is curved; it's a parabola, see Figure 1.21. Suppose you want to know how the function behaves in the neighbourhood of a certain point, say at  $x = 1$ . If you zoom in at the graph of  $x^2$  around that point, you'll notice that the graph near  $x = 1$  looks remarkably like a line, see also Figure 1.21. Of course, the graph still curves a little bit, but you don't notice that when you're zoomed in that much. Therefore, if you stay close enough to  $x = 1$ , you can approximate the graph of  $x^2$  very well by a line. This line is called the tangent line, and its slope is exactly equal to the slope of the graph of  $x^2$  at  $x = 1$ . Because you can calculate the slope of a graph of a function by calculating its derivative at that point, you see that taking the derivative of a function has everything to do with obtaining 'local' information of that function, near a certain point. Of course, in our example, the choice of  $x = 1$  as a point around which we wanted to obtain local information was completely arbitrary, we could have chosen any other point and taken a close look at the graph around that point. For that matter, we could have chosen a different function to investigate: the principle of obtaining local information by using the derivative stays the same.

How exactly do you use the derivative to obtain local information about a function around a certain point? Let's look at the example again, and let's zoom in again at  $x = 1$ . The slope of the graph at that point is 2: if you start at  $x = 1$  and move a tiny bit to the right, the function value will increase twice that tiny bit, see Figure 1.22. If we call that tiny increment  $\delta$ , we see that the tangent line around  $x = 1$  is given by

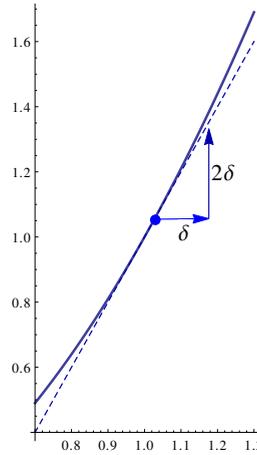


Figure 1.22: The linear approximation of  $x^2$  near  $x = 1$ .

$1 + 2\delta$ . Therefore, if  $\delta$  is small, the real function value  $(1 + \delta)^2$  is approximated very well by the value at the tangent line,  $1 + 2\delta$ . In other words,

$$(1 + \delta)^2 \approx 1 + 2\delta \text{ if } \delta \text{ is small.}$$

This is called a *linear approximation* of the function  $x^2$  near  $x = 1$ . You can approximate any function at a certain point in this way, by using the slope of the function at that point. So, in general, if you have a certain function  $f(x)$  which you want to approximate at a certain point, say  $x = 1$ , the linear approximation gives

$$f(1 + \delta) \approx f(1) + \left. \frac{df}{dx} \right|_{x=1} \times \delta$$

where  $\left. \frac{df}{dx} \right|_{x=1}$  denotes the derivative of the function  $f$ , calculated at the point  $x = 1$ , i.e. the slope of the graph of the function  $f$  at  $x = 1$ .

Now it's time to make a conceptual jump. Imagine that the point  $x = 1$  represents the pulse, and the function  $f(x)$  represents the reaction-diffusion system. The idea is to use this principle of linear approximation to find out information about the behaviour of the reaction-diffusion system 'near' our pulse. Of course, it will not be possible to draw this anymore, or to visualise it properly in another way. However, the approach is exactly the same as in the previous example.

We take our pulse, and add a little disturbance to it. Since the pulse is a stationary

solution to the reaction-diffusion system, it will not change in time. However, the pulse with a little added disturbance *will* in general change in time if we look how it evolves according to our reaction-diffusion system. As explained above, we want to know if the little disturbance ‘dies out’, i.e. diminishes in time, or if it will grow in time. The behaviour of the little disturbance characterises the stability of our pulse. This disturbance has to be rather general, since we don’t want the stability of our pulse to depend on the random choice of our disturbance. Maybe if you disturb the pulse a certain way, the disturbance will diminish, while if you disturb the pulse in a slightly different way, that disturbance will grow in time. If your pulse is to be stable, you want *all* disturbances to diminish in time, since in practice, you don’t have control over the way your pulse is disturbed. Even if there’s just one specific way to nudge the pulse such that that disturbance grows in time, you cannot say that the pulse is stable, so then the pulse is unstable.

Back to the pulse and its disturbance. The pulse is a stationary solution to our reaction-diffusion system, so it has a  $U$ -component and a  $V$ -component which are both independent of time but still depend on the spatial variable  $x$ . Let’s call these pulse components  $U_{\text{pulse}}(x)$  and  $V_{\text{pulse}}(x)$ . If we disturb the pulse, we disturb both components, maybe in different ways. Let’s call the disturbance in the  $U$ -component ‘ $u$ ’ and the disturbance in the  $V$ -component ‘ $v$ ’. Note that these lower case  $u$  and  $v$  have nothing to do with the same symbols which were introduced in section 1.3.1, in the context of the existence question. Although it can be a little bit confusing to re-use symbols in this way, it’s often more convenient than introducing yet another symbol.

If we want to see whether the disturbances  $u$  and  $v$  grow or diminish in time under the evolution described by the reaction-diffusion system, we’ll have to put them in there and see what happens. In other words, we make the substitution

$$U \rightarrow U_{\text{pulse}} + u, \quad V \rightarrow V_{\text{pulse}} + v$$

Here, the pair  $(U_{\text{pulse}} + u, V_{\text{pulse}} + v)$  plays the role of ‘ $1 + \delta$ ’ in the previous example. By the principle of linear approximation, we can use the ‘derivative’ of the reaction-diffusion system calculated ‘at the pulse  $(U_{\text{pulse}}, V_{\text{pulse}})$ ’. It’s too technical to go into details at this point, but the result is that we obtain two linear equations: one for  $u$ , one for  $v$ . Although they are still intertwined (the equation for  $u$  depends on  $v$  and vice versa), they are intertwined in a simple, linear way.

For such systems of linear equations, it’s possible to isolate the growth rate of  $u$  and  $v$ . Remember,  $u$  and  $v$  were still very general disturbances, which are functions of the spatial variable  $x$  (since they deform the pulse) and the temporal variable  $t$  (since

they evolve in time). This growth rate is called  $\lambda$ , and its value determines the fate of the disturbances – both disturbances have the same growth rate, that’s one of the advantages of having a set of linear equations. The way the growth rate determines the stability is very straightforward: if  $\lambda$  is negative (or, more specifically, if its real part is negative, but that’s not really important here), the disturbances  $u$  and  $v$  will decay, so then the pulse is stable. If  $\lambda$  is positive, the disturbances will grow, so the pulse is then unstable.

If you introduce this growth rate  $\lambda$  and you start to shove terms around a little bit, you end up with a system of four linear equations, in this case a four-dimensional linear dynamical system. It’s rather similar to the way we obtained a four-dimensional dynamical system in section 1.3.1: the four components are given by  $u$ , the derivative of  $u$ ,  $v$ , and the derivative of  $v$ . That’s just like in section 1.3.1: even the symbols are the same, even though they represent something else. Also, this dynamical system is linear, whereas the dynamical system in section 1.3.1 was not linear. Does that make things easier? Yes and no. Although our newly obtained linear dynamical system is linear, it has a peculiarity which makes it in principle hard to analyse: it depends explicitly on  $x$ . We can write the system down in a concise way as

$$\frac{d}{dx}\phi = A(x; \lambda)\phi$$

where we put all four components  $u$ ,  $\frac{d}{dx}u$ ,  $v$  and  $\frac{d}{dx}v$  into one four-component vector  $\phi$ . As you can clearly see, this is an example of a dynamical system, only in terms of  $x$  instead of  $t$  (see section 1.2.2). From a mathematicians’ point of view, this makes absolutely no difference: it’s just a symbol. This is an example of a situation where it’s useful to ‘forget’ for a moment what  $x$  stands for, allowing you to recognise the system of equations for  $u$ ,  $v$  and their derivatives as a dynamical system. This allows you to use techniques from the field of dynamical systems in a situation where you initially wouldn’t have thought they could come in handy. The symbolic language of mathematics shows an unforeseen connection between this analysis of small disturbances and dynamical systems.

I mentioned that this dynamical system we’ve cooked up is hard to analyse since the system depends on  $x$  – in particular, the matrix  $A(x; \lambda)$  (since that’s what it is, a matrix) depends on  $x$ . Don’t worry about matrices and vectors: the important thing is that this dynamical system looks different for different values of  $x$ . Why is this a problem? Well, remember that the variable  $x$  has taken over the role of ‘time’ in this dynamical system. A dynamical system which depends actively on time is constantly changing the evolution rules as you let an initial state evolve in time. It’s like rolling

a ball over a table top, while constantly wiggling and moving the table itself. You can imagine that the path of the ball becomes very erratic, and will be highly influenced by the movements of the table.

Luckily, in our case, we have some knowledge of how  $A$  depends on  $x$ , in other words, we know something about the ‘movements of the table’. The reason  $A$  depends on  $x$  is that we consider small disturbances  $u$  and  $v$  of our pulse solution, i.e. we’ve started to look ‘nearby’ our pulse. Therefore, we had to calculate the ‘derivative’ of the reaction-diffusion system ‘at the pulse’. Regardless of its exact meaning, this ‘derivative’ of the reaction-diffusion system evaluated ‘at the pulse’ is exactly what defines  $A$ . Therefore, the spatial dependence of  $A$  is very closely related to the spatial structure of the pulse.

Remember, our pulse solution looks like the one in Figure 1.16. It has three important features:

- a) The  $V$ -component of the pulse is very sharp and very narrow: except for a very small region in the middle, the  $V$ -component is practically flat and almost zero.
- b) The  $U$ -component of the pulse is, except for the very small region in the middle, not really influenced by the  $V$ -component, since the  $V$ -component vanishes almost everywhere.
- c) Both components decrease and become very flat and almost zero as you move further and further away from the pulse peak, i.e. if  $x$  becomes very large or very negative.

This spatial structure of the pulse has a direct influence on the way  $A$  depends on  $x$ . In the small central region, it is predominantly determined by the  $V$ -component of the pulse, whereas outside that region, the  $U$ -component of the pulse takes its role in determining the spatial dependence of  $A$ . Since both components become flat and almost zero when you are far away from the pulse peak,  $A$  loses its  $x$ -dependence and becomes very simple in that far-away limit. In Figure 1.23, the  $x$ -dependence of  $A$  is visualised based on the spatial structure of the pulse.

With all this talk about the dynamical system and  $A$ , it seems that we’ve drifted away from our original goal, which was to determine the growth rate  $\lambda$  of our disturbances  $u$  and  $v$ . How do we combine these two things?

First, it’s good to notice that our dynamical system is an equation for  $u$  and  $v$ , and the dynamical system uses the variable  $x$ . Therefore, the dynamical system determines the spatial dependence of  $u$  and  $v$ , in other words, the way they deform the pulse. So,

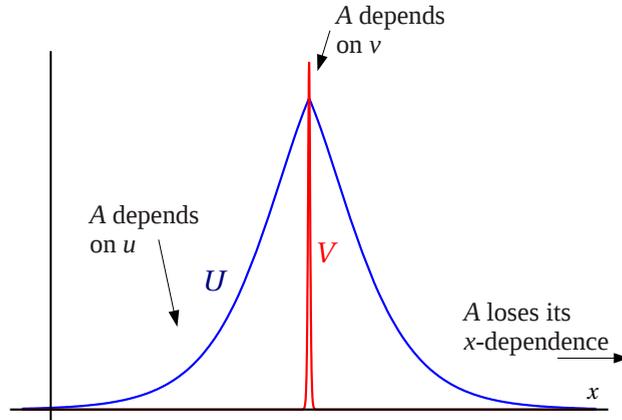


Figure 1.23: The  $x$ -dependence of  $A$ , based on the spatial structure of the pulse.

we're looking for small disturbances  $u$  and  $v$  which are at the same time solutions to the dynamical system  $\frac{d}{dx}\phi = A(x; \lambda)\phi$ . Are all possible solutions of this dynamical system then disturbances? Certainly not. The reason for this is that disturbances are, by nature, small – that means, they are small *everywhere*, for all values of  $x$ . However, not every solution of the dynamical system stays small for all  $x$ . It's even worse: in general, a solution to such a dynamical system is 'unbounded'. That means that such a solution will increase as  $x$  increases, and it can become as large as you want if you go to higher and higher values of  $x$ . Analogously, things can go wrong at the other side: a solution can become as large as you want if you go to more and more negative values of  $x$ . A good example is the function  $x^2$ , see Figure 1.21: although it's fairly small near  $x = 0$ , it increases as  $x$  becomes larger or more negative, and this never stops. You can obtain arbitrarily high values if you just go far away enough from  $x = 0$ . Therefore,  $x^2$  is unbounded, and that's the reason that it is not admissible as a small disturbance.

You can wonder if it's even possible at all to find bounded solutions for the dynamical systems, which are then admissible as small disturbances. The answer is yes, well, sometimes: it depends on  $A(x; \lambda)$ . Now we remember that  $A$  also depends on the growth parameter  $\lambda$ . Here,  $\lambda$  plays the role of a parameter: for different values of  $\lambda$ , the matrix  $A(x; \lambda)$  will be different, and therefore the associated dynamical system will be different, and the possible solutions to that dynamical system will be different for different values of  $\lambda$ . You can imagine that for some values of  $\lambda$ , the dy-

dynamical system doesn't have any bounded solutions, while for other values of  $\lambda$  there are bounded solutions possible. Finding bounded solutions, and therefore obtaining admissible disturbances, is now a question of choosing the right values of  $\lambda$ .

Those 'right' values of  $\lambda$  are called 'eigenvalues', and since they were originally introduced as growth rates, they determine the stability of the pulse. The strategy to determine the stability of the pulse is therefore as follows: first, we determine for which values of  $\lambda$  the associated dynamical system admits bounded solutions. Then, we take a look at those eigenvalues: if all of them are negative, then the pulse must be stable. This is because every admissible disturbance is, since it is a solution to the dynamical system, associated to some eigenvalue. This eigenvalue is the growth rate of that disturbance. If you know all eigenvalues, you know the growth rates of all admissible disturbances. If all disturbances have a negative growth rate, then they will all diminish, and therefore the pulse is stable. However, if there is just one positive eigenvalue, the disturbance (or disturbances) associated to that eigenvalue will grow. Therefore, not every disturbance will diminish in time, and therefore the pulse is unstable.

In general, it's very hard to obtain such eigenvalues, let alone obtain them all. If we want to know whether a pulse is stable, we need to know all eigenvalues: we want to make sure that each and every one of them is negative, or else the pulse is not stable. On the other hand, once we found just one positive eigenvalue, we can immediately conclude that the pulse is unstable. Therefore, it's often much easier to prove statements like 'Under these and these circumstances, the pulse is unstable' than statements like 'Under these and these circumstances, the pulse is stable'. In chapters 2 and 3, you will therefore find more instability results than stability results. However, in some cases, it is possible to prove stability, especially in chapter 2 (section 2.4.2).

Our knowledge of the  $x$ -dependence of  $A$  can help us to obtain these eigenvalues. What's really helpful, is that we know that  $A$  becomes very simple in the far-away limit, if we take  $x$  to be very large or very negative. From this very simple 'limit'-version of  $A$ , which doesn't depend on  $x$  anymore, we can infer some properties of the solutions of the entire dynamical system. Since  $A$  becomes very simple as  $x$  becomes very large (or very negative), we know that the solutions of the dynamical systems should behave accordingly. From this 'limit'-version of  $A$ , there are just two possibilities: a solution either grows in an unbounded way as  $x$  becomes larger and larger, or it decays to zero as  $x$  becomes larger and larger. Therefore, we reason as follows. Whatever a solution to the dynamical system looks like 'in the middle', where  $A$  depends on some complicated way on  $x$ , if we want that solution to be bounded, i.e. to

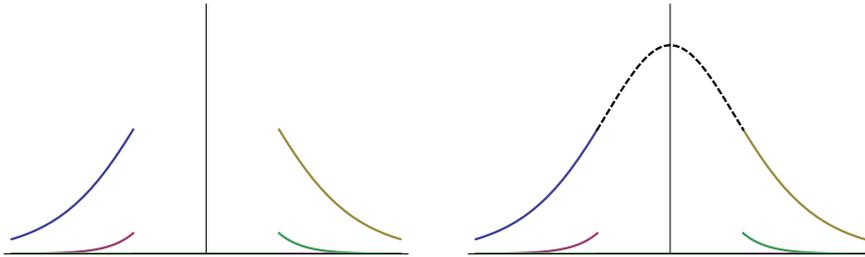


Figure 1.24: Matching solutions which decay to the left and to the right.

be an admissible disturbance, it has to decay to zero as  $x$  becomes very large or very negative. Because an admissible disturbance, i.e. a bounded solution has to decay to zero both as  $x$  becomes very large and as  $x$  becomes very negative, we try something similar as in section 1.3.1, where we tried to construct a pulse: we approach it from both sides. That is, we take a solution which decays as  $x$  becomes very negative (which behaves nicely at ‘the left’); we take a solution which decays as  $x$  becomes very large (which behaves nicely at ‘the right’), and we try to let them match. In Figure 1.24, this idea is illustrated.

In section 1.3.1, we tried to let two sheets intersect, see Figure 1.15. This ‘intersection criterion’ could be calculated, and it gave the value which the ‘constant’ component  $u$  should have in the very small zone where the sharp  $v$ -spike was. In other words, it gave the value of that parameter, such that the sheets would intersect. While matching these decaying solutions from both sides, something equivalent is going on. When trying to tune the value of  $\lambda$  such that these solutions can indeed be matched, you obtain a ‘matching criterion’, which gives you the ‘right’ values of  $\lambda$ , in other words, the eigenvalues. This ‘matching criterion’ can be reformulated in terms of a function, which is called the Evans function. This Evans function is particularly useful since the values for which the Evans function is zero are precisely the eigenvalues you’re looking for. In short: if you know the Evans function, you know the eigenvalues, and therefore you know whether the pulse is stable or not. It turns out that the scale separation which is present in our pulse, and which had a direct influence on the  $x$ -dependence of  $A$  (see Figure 1.23), can be used to separate the equations for the disturbance  $u$  and the disturbance  $v$ , just as in the existence section 1.3.1. Analysing those equations separately, it is possible to obtain the Evans function explicitly in the end. It’s too much to go into details, but the similarities between the approach of the existence problem and the stability problem are striking. In the

end, the process of ‘matching’ solutions to obtain admissible disturbances is just the problem of letting two sheets of properly chosen solutions intersect. Behind it all, there’s the singularly perturbed nature of the reaction-diffusion system which causes the scale separation, which in turn allows us break up the large problem into smaller sub-problems. Then, you analyse these sub-problems separately and combine the results of these analyses at the end. Without the singular perturbation, none of this would be possible.

### 1.3.3 Dynamics

In the previous section, I’ve explained how you can analyse the stability of the pulse whose existence was established in section 1.3.1. It was already mentioned in section 1.3 that reaction-diffusion systems often depend on a number of parameters, which you can choose freely. Depending on the value of these parameters, a pulse might exist, and if it exists, it can be stable; see Figure 1.11. The ‘existence zone’ in parameter space can be found by the methods explained in section 1.3.1; the ‘stability zone’ inside the existence zone can be found by the methods explained in section 1.3.2.

In this section, I’ll elaborate some more on the fate of the pulse when it is disturbed. If the pulse is stable, it’s easy to see what will happen: the disturbance will fade away, and the pulse will stay where it is, in its original shape. However, if the pulse is unstable, it’s not clear at all what will happen. Yes, the disturbance will grow, but how does that affect the pulse? Will the pulse grow as well? Will it deform? Will it start to move? In general, these questions are very hard to answer.

The easiest way to get an idea about what’s going to happen with an unstable pulse is to do numerical simulations. You put your reaction-diffusion system in encoded form in a simulation program, you prescribe your pulse, disturb it a little, push the button, and see what happens. Since the computer does all the work, this seems like a very good approach – although the actual setup of such a simulation can be a lot harder than you might imagine. However, this approach has its limitations. You have to choose a value for each of your parameters, including  $\varepsilon$ ; since the computer calculates things step by step, you have to tell the computer to take very small time steps – but not too small, otherwise the simulation takes too much time. Also, you have to divide the space into small parts, because the computer can only handle numbers, and not smooth things like functions. In the case of our pulse, where the  $V$ -component is very narrow and very sharp, you’ll have to divide the space in that region into a large number of very, very small segments.

All these things introduce a certain aspect of trial and error. You have to make sure

that  $\varepsilon$  is ‘small enough’, whatever that may mean, and you will have to be lucky in picking the values of your parameters such that something interesting happens. If you’ve got a lot of parameters, this becomes increasingly difficult. Moreover, if something exciting happens, you often don’t really know *why* this happens. Therefore, numerical simulations can act as a guide to interesting phenomena, but the real insight often has to come from ‘proper’ mathematical analysis.

There are situations where you can analytically investigate the fate of an unstable pulse. This is when a pulse *becomes* unstable: by that, I mean the following. Suppose you know that for a certain parameter choice, your pulse is stable, i.e. you’re in the ‘stable region’ in parameter space, see Figure 1.11. Then, you start to change the value of one of the parameters, and see what will happen to the growth rates, the eigenvalues. At some point, one of your eigenvalues (which were all negative) will become positive. In other words, you’ll cross the boundary of the ‘stable region’ in Figure 1.11. At the boundary of that region, where the pulse becomes unstable, it’s possible to see what’s happening in more detail, i.e. to analyse *how* the pulse becomes unstable.

Such a crossing of the boundary of a region in parameter space is an example of a ‘bifurcation’. A ‘bifurcation’ is a rather general term to describe a situation where something genuinely changes. That’s what happening at the boundary of the stable region: within that region, you can change your parameters around a little, but that doesn’t have large effects: the pulse will change its shape a little, but it will still be stable. However, if you cross the stability boundary, this really makes a difference. If you pick two parameter values which are very close together but on opposite sides of the stability boundary, the two associated pulses will look very much alike; however, one pulse will be stable, while the other pulse is unstable. Something similar happens when you cross the boundary of the existence region: suddenly, the pulse that you had can no longer exist.

Our pulse can become unstable in a number of ways. For each way of becoming unstable, there is an associated bifurcation, where a parameter crosses the boundary of the stability region. It really depends where you cross the boundary: if you’re unlucky, you’ll be not only out of the stability region, but immediately out of the existence region as well – see Figure 1.11. Other parts of the stability boundary are a lot less dangerous. For our kind of pulses, it turns out that there is a certain general way in which such a pulse loses its stability, and that’s through a so-called oscillatory instability. It’s not to say that this will *always* happen when you start playing with parameters and push the pulse out of the stability region, but it will happen most of

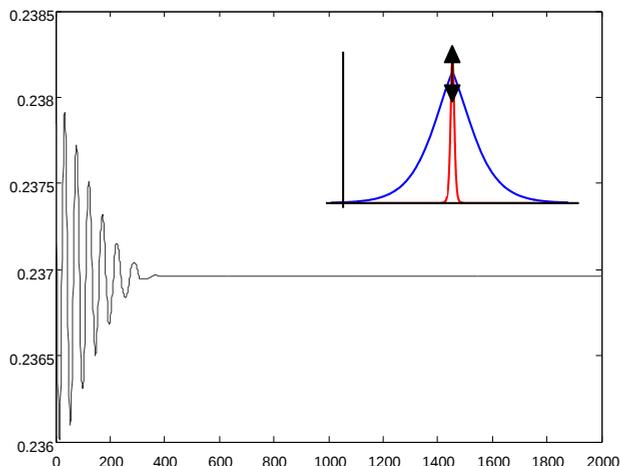


Figure 1.25: The movement of the tip of the pulse just before a Hopf bifurcation.

the time. It's even stronger: you'll have to choose your parameters very carefully to avoid this oscillatory instability while pushing the pulse out of the stability region. This oscillatory instability bifurcation is also called a Hopf bifurcation, and it manifests itself as follows. If you've got a stable pulse and you're near the boundary of the stable region, and you give your pulse a little nudge, then it will start to oscillate, i.e. to wobble up and down. This oscillation will die out after some time, because the pulse is stable. In Figure 1.25, you can see the results of a numerical simulation of such a pulse near the stability boundary, i.e. near a Hopf bifurcation. Once you change your parameters such that you cross the stability boundary and then give your pulse again a little nudge, you'll see that the pulse will start oscillating again. Only now, because the pulse is unstable, the amplitude of the oscillation will grow, and the pulse will go up and down in an increasing way. This oscillation can become so wild that the pulse can slap itself flat, and disappear. This is indeed what you observe in simulations, see Figure 1.26.

However, when playing around with parameter values and simulating the resulting pulse, you will stumble upon something quite exciting. In Figure 1.27, you can see the result of such a simulation. What you see is a pulse which is disturbed and starts to oscillate. The amplitude of the oscillation starts to grow – nothing new so far, it means that we're outside the stability region. But then, after a while, you ob-

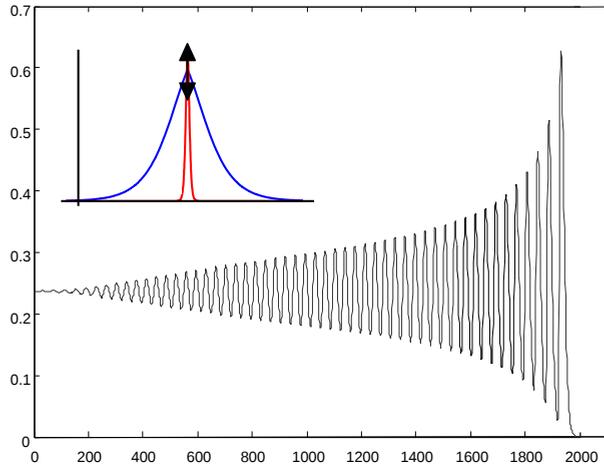


Figure 1.26: The movement of the tip of the pulse just after a Hopf bifurcation.

serve that the amplitude of the oscillating pulse stabilises, and that you're looking at a steadily oscillating, or 'breathing' pulse. This is quite something: we've discovered a pulse which moves up and down in time, a first step beyond the realm of stationary patterns. Moreover, the pulse seems to stay where it is: it wants to move up and down, but it apparently doesn't want to move to the side.

Those observations are a reason for further analysis, and a few questions arise: 'How can we go beyond the Hopf instability bifurcation and find a breathing pulse?' 'Can we predict when we will see a breathing pulse, and when we will just see an increasingly wild oscillation?'

These questions were the driving force behind the last, fourth chapter of this thesis. This chapter is also rather technical; however, it certainly is possible to give you an idea about the approach you can take to answer these questions.

In order to do that, we'll revisit the example concerning the linear approximation, section 1.3.2. There, I argued that when you zoomed in on the graph of a function, you could approximate that graph very well by its tangent line. That was called a linear approximation, see Figure 1.22. If you zoom out a little, you'll see that this linear approximation starts to differ from the 'real' graph of the function. Therefore, if you want to approximate this graph a little better, you will have to go beyond this linear approximation. This means you'll have to make a quadratic approximation.

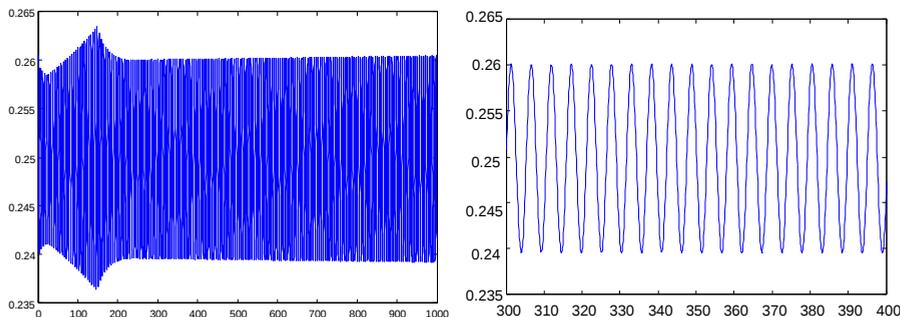


Figure 1.27: A breathing pulse. After a period of initial amplitude growth, the oscillation of the tip of the pulse settles down to a steady breathing motion, with constant amplitude. The right figure is a zoom of the left figure; here you can clearly see the periodic motion.

Formula-wise, this means that you approximate the function at a certain point (say, again,  $x = 1$ ) as follows:

$$f(1 + \delta) \approx f(1) + \left. \frac{df}{dx} \right|_{x=1} \times \delta + \frac{1}{2} \left. \frac{d^2f}{dx^2} \right|_{x=1} \times \delta^2$$

If you compare this with the linear approximation in section 1.3.2, you'll see that there's just an extra term. In that extra term,  $\left. \frac{d^2f}{dx^2} \right|_{x=1}$  denotes the second derivative of  $f$ , calculated at  $x = 1$ .

Such a higher order approximation is also possible in the case of the pulse and the reaction-diffusion system. Just like in section 1.3.2, you have to make the conceptual jump where you imagine that the point  $x = 1$  represents the pulse, while the function  $f(x)$  represents the reaction-diffusion system. Just as it is possible to take the 'derivative' of this reaction-diffusion system, you can also take its 'second derivative'. It's all quite analogous, really. The only disadvantage of this higher order approximation is that your formulas become a lot, lot longer; you can see that when browsing through chapter 4. Therefore, it also becomes harder and harder to say something definite about the fate of the pulse, because your equations become increasingly complicated. Once again, the singular perturbation helps you out. Based on the results of the existence and stability of the pulse, it turns out that it is possible to obtain explicit formulas which tell you something about whether such a breathing pulse will or will not appear.

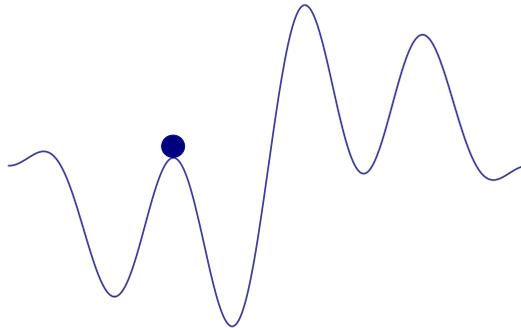


Figure 1.28: An example of a ‘landscape’.

The reason why this higher order approximation can shed light on the appearance of a breathing pulse can be made clear by an example. In section 1.3.2, the concept of stability was illustrated by a ball on top of a hill versus a ball at the bottom of a pit. The first situation was unstable, the second was stable; see also Figure 1.20. In this case, the linear approximation only tells you if you’re on top of a hill, or at the bottom of a pit. It doesn’t tell you anything of what the rest of your surroundings look like. There might be other hills nearby, or other pits: see Figure 1.28. The higher order approximation lets you take a better look at the surrounding landscape.

In the case of a breathing pulse, the following situation occurs. Since you’re in an unstable situation, you know that your ball (which represents the pulse) is on top of a hill. If the ball is given a little nudge, it will roll down. However, in some situations, your starting hill lies in between to larger hills. You might be able to roll down the hill, but you’ll never escape the larger valley: see Figure 1.29. That means that, while the ball rolls around in the larger valley, it will never escape. The associated periodic motion of the ball is therefore bounded: its amplitude doesn’t grow indefinitely.

Therefore, if you use higher order approximations and find out that the local ‘landscape’ is such as in Figure 1.29, then you know there has to exist a steadily oscillating, breathing pulse. Moreover, you can find that particular breathing pulse by giving an unstable stationary pulse nearby a little nudge, such that it starts to oscillate (see for example Figure 3.1). In chapter 4, the analysis tells us that such a situation can indeed arise, if you choose your parameters correctly. In more technical terms, that means that the Hopf bifurcation is called ‘supercritical’. If there is no surrounding valley, and your ball keeps rolling down the hill with increasing speed, then the Hopf bifurcation is called ‘subcritical’. Using these concepts, it is possible to prove the existence of a breathing pulse.

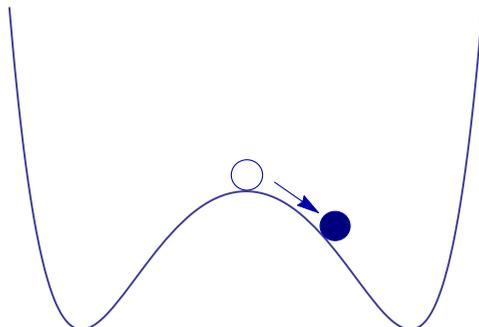


Figure 1.29: A hill in a valley: bounded periodic motion.

## 1.4 Thesis structure and main results

Now that you've encountered some central concepts and methods used in this thesis, it is time to explain how this thesis is structured, and which results have been obtained during the research leading to this thesis.

Let's focus first on chapters 2 and 3. I'll introduce them together, since they are rather similar in structure and content.

Chapter 2 treats pulse solutions in the 'slowly nonlinear Gierer-Meinhardt model'. The Gierer-Meinhardt model, mentioned in section 1.2.2, is one among many canonical models which have been studied in the last few decades in the context of pattern formation in reaction-diffusion systems. While the origin of the Gierer-Meinhardt model lies in developmental biology (it was developed in the context of morphogenesis), other sciences have contributed to the class of canonical reaction-diffusion models as well. The Gray-Scott model [23] has its basis in chemistry, while the Fitzhugh-Nagumo model [20, 40] describes the transmission of a electrical signal in a nerve.

The study of patterns in reaction-diffusion equations has largely been centered around the models just mentioned. One of the distinguishing features of the research presented in this thesis is the step beyond these canonical, quite specific models. Once you know how to find a pattern in a reaction-diffusion system, it doesn't matter that much what the system looks like. Of course, there are some mild restrictions on the reaction terms in the reaction-diffusion system under consideration, but those aren't really 'restrictive'; still, a wide range of systems falls into this admissible class.

The slowly nonlinear Gierer-Meinhardt model is such a step beyond the canonical models, in this case beyond the Gierer-Meinhardt model. In chapter 2, the existence and stability (in that order) of a pulse in this extended version of the Gierer-Meinhardt model is treated. It turns out that the extension of the Gierer-Meinhardt model with a slowly nonlinear term introduces a lot of new things: not only is it possible to construct a new, different type of pulse in comparison with the ‘canonical’ Gierer-Meinhardt model, the stability analysis of the constructed pulse shows that it can become unstable in a previously unobserved way. Another feature of this model is that, even though it is a nonlinear extension of the canonical model, you can still calculate everything you want in an explicit way: there is an explicit formula for the pulse, the stability analysis can be carried out explicitly, and even the Evans function can be determined in an explicit way. This is very useful for the stability analysis, in finding all the eigenvalues. For this model, it is possible to prove some results about the stability of the pulse.

Chapter 3 has the same structure as chapter 2. The main difference is that here we treat a wide range of possible reaction-diffusion systems instead of just one specific example. Since we showed in chapter 2 that such an extended system can give new results, it’s interesting to see how far you can go with this. During the entire analysis in chapter 3, the reaction-diffusion system is kept at a very general form, almost like ‘something involving  $\phi$ ’ as in section 1.2.2 – well, it’s a little more exact than that, but the main idea still is that it doesn’t really matter that much what your reaction-diffusion system looks like. Once it obeys some very general conditions, it is possible to construct a pulse solution for it. Even the stability analysis can be carried out in the same way as for the explicit example of the slowly nonlinear Gierer-Meinhardt system of chapter 2. What’s even more surprising, is that for such a very general system, you can still write down the Evans function explicitly. This is a really strong result, because it allows you to prove a number of (in)stability results for this very general class of systems. It even becomes clear that already during the construction process of the pulse, you can predict which pulse will certainly be unstable, and which pulse has a chance to be stable. The previous sentence already suggests that there’s more than one possibility to construct a pulse (just like in chapter 2), which is indeed true in general. Chapter 3 shows you how to approach such a general system, and gives you the tools to construct and analyse a pulse yourself.

In this sense, chapters 2 and 3 are each other’s mirror image. They follow exactly the same approach: chapter 2 for an explicit system, chapter 3 for a very wide class of systems of which the slowly nonlinear Gierer-Meinhardt system of chapter 2 is just an example.

While both chapters 2 and 3 treat existence and stability, the first two steps of the general research approach (Figure 1.10), the last chapter takes a first step into the largely unexplored world of pulse dynamics. Here, the higher order approximation of a destabilising Hopf bifurcation (which was found to exist in general in chapter 3, and was explicitly established in chapter 2) leads to the discovery of so-called breathing pulses, see section 1.3.3. In this fourth chapter, I present two equivalent methods to calculate these higher order approximations. These higher order approximations are applied to the slowly nonlinear Gierer-Meinhardt system, using the results of chapter 2. The numerical simulations which were presented at the end of chapter 2, section 2.5 already hinted at the existence of breathing pulses (see also Figure 1.27): in chapter 4, it is shown that these solutions can indeed exist, and when this is the case.