Nature's Patterns: a tapestry in three parts

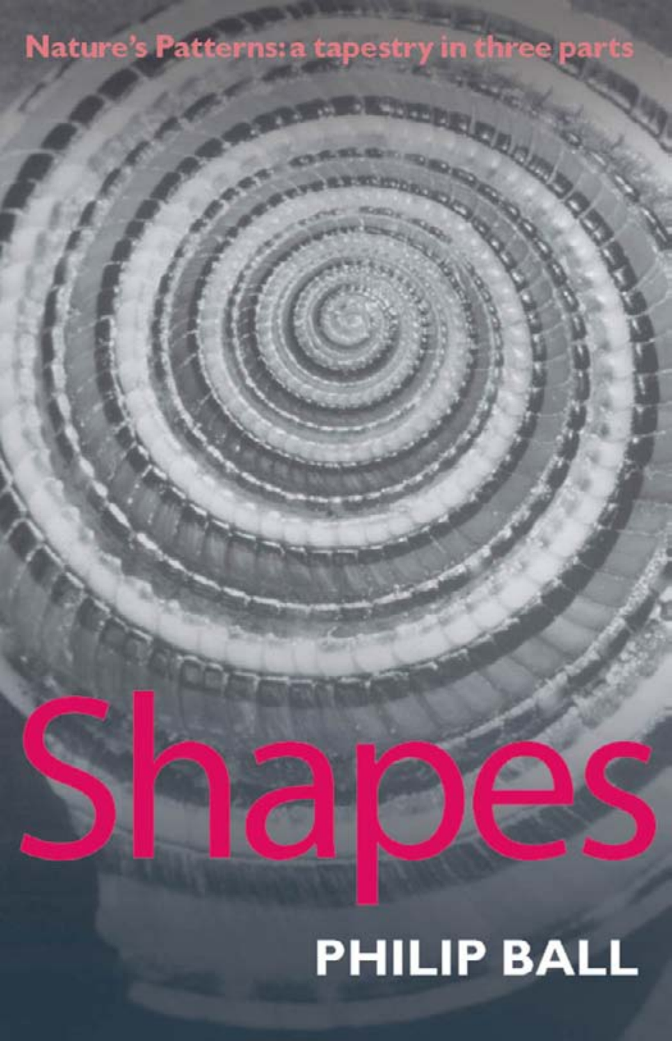

# Shapes

PHILIP BALL

(called the latent heat) which warms the surroundings, enabling it to ‘afford’ the order of the crystalline state.

The key fact here is that, one way or another, an equilibrium state exists for any system. And the second law of thermodynamics dictates that this is the state for which the total entropy of the universe is maximized. Moreover, the second law seems to insist that spontaneous change will always carry the system towards this equilibrium state. This change might take a very long time—iron can take decades to turn into rust—but the direction is always clear. Thus, the second law appears to define an arrow of time: change happens in the direction that takes the universe from a low-entropy to a high-entropy state.

And now we come to the second of the two possible pattern-forming processes that I alluded to above. These are ones in which the direction of change does *not* lead us towards an equilibrium state corresponding to the state of highest entropy. Such processes, indeed, sometimes appear to be leading in the opposite direction. This looks at first glance as though it violates the second law of thermodynamics, the very principle that I have just asserted to be universal. That is exactly what scientists in the mid-twentieth century thought too, and so they were most disapproving when Boris Pavlovitch Belousov claimed to have evidence to the contrary.

## Out of balance

Belousov, a Russian biochemist, was not looking for controversy. By all accounts he would have been glad of a quiet life; but in the 1950s he discovered something that he could not ignore, no matter how heretical it seemed.

He was interested in the metabolic process called glycolysis, by which enzymes break down glucose and capture the energy that this chemical reaction releases. Belousov devised a cocktail of chemical ingredients that was supposed to represent an artificial analogue of glycolysis, and he mixed them together. But the reaction that followed did not seem to settle down into an equilibrium state. The mixture was initially clear, and it turned yellow as the reaction proceeded. But having done so, it then turned clear again, and then back to yellow and so on, pulsing at regular intervals over and over again.

Chemical reactions, like all other processes of change, have a ‘downhill’ direction: that which leads to an increase in total entropy. When the reaction proceeds, the second law seems to insist that it must do so in this direction: the reaction eventually reaches the equilibrium state in which the total entropy change has been maximized.* But Belousov seemed to be suggesting that his reaction had no preferred direction: first it went one way, then the other. It was as though he was claiming to observe an ink drop that dispersed and then reformed, time and time again.

That sounded absurd, and so Belousov found himself unable to publish his findings in any reputable journal. Everyone decided that they were obviously due to his experimental incompetence. In the end, he sneaked the results into an obscure volume of conference proceedings on a completely different topic. Outside the Soviet Union, Belousov’s ‘oscillating reaction’ remained unknown.

The irony was that Belousov’s discovery was not new, and neither did it lack an explanation. In 1910 the Austrian-American ecologist and mathematician Alfred Lotka described how in theory a chemical reaction might undergo just this kind of oscillation, switching its direction back and forth. In his original model, the oscillations were ‘damped’—like the decay of a ringing bell, they gradually die out and the system settles into a steady state. But ten years later Lotka showed how such oscillations might be sustained indefinitely.

Lotka was not one to defer to the imperatives of thermodynamics. ‘The two fundamental laws of thermodynamics’, he wrote, ‘are, of course, insufficient to determine the course of events in a physical system. They tell us that certain things cannot happen, but they do not tell us what does happen.’ How could he make such a claim? Lotka had the insight, obvious now but remarkable at the time, that living systems (and not just living systems) are unlike a flask of chemicals mixed together and left to react. They are, in contrast, constantly

---

*At the level of individual molecules, chemical reactions *do* run in both directions. If the reaction involves molecule A joining to molecule B, then, were we to be able to follow the process through a microscope that gave us a molecular-scale picture, we would see both A and B combining *and* the composite AB molecules falling apart, even at equilibrium. But at equilibrium the rates of these two processes are equal, so that the average amounts of A, B and AB stay the same. Depending on the magnitude of the overall entropy change, this equilibrium state may contain a greater or lesser proportion of AB. This is an illustration of the fact that the molecular world is always dynamic, its molecules always in motion. The stasis of equilibrium is a reflection of the unchanging averages that emerge from this microscopic dynamism.

acquiring energy from their surroundings: plants soak up sunlight to conduct photosynthesis, and organisms ranging from bacteria to humans gobble up energy-rich matter such as plants in order to drive their metabolic biochemical processes. '[In] systems receiving a steady supply of available energy (such as the earth illuminated by the sun)', said Lotka, 'and evolving, not towards a true equilibrium, but (probably) toward a stationary state, the laws of thermodynamics are no longer sufficient to determine the end state.' In other words, a constant flux of energy can prevent equilibrium from being reached. It is significant that Lotka says this applies not only to living systems but to our planet as a whole, bathed in the energy of the sun. We will find these prescient words, written in 1922, echoing throughout the rest of this book.

D'Arcy Thompson mentions Lotka's work in the revised edition of *On Growth and Form*, but he did not appreciate its full implications. It crops up not in any discussion of chemistry or biochemistry but in Thompson's description of animal population dynamics. For this was Lotka's principal focus; his scheme for an oscillating chemical reaction was proposed primarily as an analogue of the way populations of animals interact, as though they are no more than molecules: 'When the beast of prey A sights its quarry B, the latter may be said to enter the field of influence of A, and, in that sense, to collide with A', Lotka explained.

Lotka formulated a series of equations that showed how oscillating population sizes might arise in these 'colliding' predators and prey. His work was extended in the 1930s by the Italian biologist Vito Volterra, who showed how the scheme could be used to understand fluctuations in fish populations. I shall return in Chapter 5 to this description of the dynamics of ecosystems.

When he presented his theory of persistent oscillations in 1920, Lotka referred in passing to the fact that 'in chemical reactions rhythmic effects have been observed experimentally'. He gave no details, and frankly I do not know to whose experiments he was referring. But such oscillations were certainly reported the following year by the chemist William Bray at the University of California at Berkeley. Bray was certainly no bungling experimenter: he pioneered the study of the rates of chemical reactions, and his student Henry Taube became a Nobel Laureate. Yet when Bray found that a chemical reaction between hydrogen peroxide and iodate seems to deliver its products, oxygen and iodine, in pulses, he got a reception as cool as Belousov's 30 years later.

Bray even cited Lotka's work as evidence that such a thing was possible, but to no avail.

It is a testament to the respect and awe with which scientists regard the second law of thermodynamics that the reality of oscillating reactions such as those proposed by Lotka took half a century to become established. During the 1960s, the biochemist Anatoly Zhabotinsky, then a graduate student at Moscow State University, came across Belousov's buried findings and decided (for graduate students tend not to have their preconceptions set in stone) to take them seriously. Zhabotinsky discovered a mixture of chemical compounds that produced a colour change far more striking than the rather insipid transformations of Belousov's solutions: the oscillations went from blue to red. This concoction, in which the organic compound malonic acid reacts with bromate, with metal atoms added to catalyse the reaction, is now known as the Belousov–Zhabotinsky (BZ) reaction. If the BZ reagents are combined and mixed well, the solution switches from red to blue and back every few minutes.

Something as dramatic as this was hard to ignore, and thanks to Zhabotinsky's insistent advocacy chemical oscillations became accepted during the 1970s as something that really happens. In 1980 Belousov and Zhabotinsky, together with their colleagues Albert Zaikin, Valentin Krinsky, and Genrikh Ivanitsky, were awarded the prestigious Lenin Prize by the Soviet government. That was not quite the happy ending it seems, for Belousov died ten years earlier, while his discovery had not yet gained wide recognition.

## The chemical see-saw

How, then, does the BZ reaction elude the second law? Well, it doesn't, and neither does any known process, whether physical, chemical or biological. If we leave the BZ reaction blinking away, we find that the oscillations do not last for ever. Eventually (it can take hours) the mixture will settle into a steady, unchanging state of equilibrium—and this is indeed a state in which the entropy of the flask and its surroundings has increased. It is simply that the mixture takes a long and circuitous route to reach this destination.

Lotka was on the right track, but he did not quite express the matter correctly. It is true that thermodynamic laws do not necessarily tell us 'what does happen', but this is because they speak only about end points. They tell us what is the equilibrium state of a system, whether that is a flask of chemicals or a planet. But they do not say anything about how that state is reached, about how the change unfolds. As chemists know, thermodynamics alone is of limited value in comprehending chemical change. It is not of much practical use, for example, to know that two reagents can be combined to make a thermodynamically stable product, if that process takes a thousand years to occur. The rate and manner with which chemical reactions take place are the subjects of the discipline known as chemical kinetics. Understanding the BZ reaction is thus a question of unravelling the kinetics of the process.

The oscillations of the BZ reaction are not unavoidably doomed to fade away. They can be sustained indefinitely if the pot is constantly replenished with the raw ingredients of the reaction, and if the end product (primarily, malonic acid to which bromine atoms have been added) is removed. Chemists have devised vessels that sustain a constant throughflow of materials, which are well mixed in a reaction chamber: they call them continuous stirred-tank reactors, or CSTRs. It is not hard to see that we too—each one of us, each living human being—are fundamentally CSTRs, filled with a chemical mixture of staggering complexity that is compartmentalized into highly intricate channels and chambers. Our genes are not what sustains our lives: they would be helpless if our internal 'stirred tanks' were not constantly replenished with raw materials (sugars, amino acids, vitamins, salts, oxygen, water) and flushed free of wastes (not just the obvious ones but also compounds such as the carbon dioxide we exhale).

This continual flux of materials through a CSTR prevents the chemical system from reaching equilibrium. You have never in your life experienced true personal equilibrium, for if you had, your life would be over. Equilibrium is deathly; nothing happens there. Equilibrium for the universe means Clausius's heat death, a cosmos rendered totally homogeneous. Scientists care about equilibrium states, but to the rest of our world they are anathema. All life exists out of equilibrium, and as Lotka observed, this is ultimately made possible on Earth by the continual flux of energy from the sun.

This is what makes the planet itself come to life. It is why water circulates between the sky and the seas, why the winds blow, why plants grow and why the biosphere sustains the atmosphere in a state of extreme disequilibrium. You cannot keep an atmosphere full of a compound as reactive as oxygen for millions of years without some non-equilibrium process to replenish it: it would otherwise react with rocks and get bound up in the solid earth. This is why scientists searching for life on other planets believe that an atmosphere with a high oxygen content is a likely signature. No need to search for cities, roads, radio broadcasts: the atmosphere alone can give the game away. As indicated in Chapter 1, some researchers have suggested that the best way to look for life on Mars is not to sift through the soil for bugs but to analyse the chemical environment for signs that it is not at equilibrium.

What Lotka showed is that a chemical system out of equilibrium can develop a kind of pattern: in this case, a pattern in time, a regular oscillation in composition. We will see shortly that non-equilibrium chemical reactions have a far more general pattern-forming potential than this. But first we need to look at where the oscillations come from. What is it that makes Lotka's hypothetical mixture, and the very real chemical cocktail devised by Belousov and Zhabotinsky, so indecisive about which direction to take?

## Blow up

Like a great many chemical reactions, the BZ reaction depends on the process known as *catalysis*. A catalyst is a substance that speeds up the rate of a chemical reaction without being changed itself. There is a nice (if messy) physical analogy for this: if you pour a fizzy drink into a glass and then add a few grains of sand, the fizzing becomes more vigorous. The sand grains act as a kind of catalyst for bubble formation. They provide sites where bubbles can form more easily, which is to say, they lower the energy needed to initiate a bubble. A chemical catalyst likewise lowers the energy needed to initiate the formation of the reaction products from the initial reagents. Most industrial chemical reactions use catalysts, since they would otherwise run too slowly to be economically viable. And almost all biochemical reactions in the body are assisted by natural catalysts, the proteins called enzymes.

What makes the BZ reaction unusual is that it makes its own catalyst. This means that one of the product molecules acts as a catalyst to speed up the formation of more product: it is self-catalytic, or *autocatalytic*. This is an example of a positive feedback process, which is self-amplifying.

Left to its own devices, autocatalysis makes the process go ever faster. A nuclear explosion involves this sort of positive feedback, as do most chemical explosions. Autocatalysis is prone to literally blowing up out of control.

But how can autocatalysis lead to oscillations, rather than simply to a runaway process? There needs to be some way of checking the propensity for blowing up. Let me illustrate what Lotka had in mind here by using his own context of predator and prey populations. Let’s say that there is a population of rabbits and a population of foxes that prey on them. Rabbits are notoriously autocatalytic: rabbits lead to more rabbits, or in other words, rabbits catalyse their own production. Given unlimited grass, a rabbit population will multiply exponentially.

Now enter the foxes. They eat rabbits; and well-fed foxes multiply too—in other words, foxes lead to more foxes, but only so long as there are rabbits around to sustain them. This is also an autocatalytic process, but it depends on the presence of rabbits. Yet foxes, no matter how well fed, die off from time to time: there is a steady rate of attrition of the fox population.* So we can write three ‘equations’ to describe how the populations of predator and prey change over time:

1. Rabbits and grass lead to more rabbits
2. Rabbits and foxes lead to more foxes
3. Foxes lead to some dead foxes

Each of these processes happens at a characteristic rate.

Suppose now that we start off the ecosystem with a few rabbits and a few foxes. The rabbits multiply quickly, and their numbers rise. This means that the foxes have an abundant source of prey. In line with step 2, the fox population also starts to rise. The trouble is that the foxes don’t know when to stop: they gorge on rabbits, and decimate the rabbit population. And then they find that there is suddenly no food, and so step 2, which relies on the presence of rabbits, can no longer take place. According to step 3, the foxes already present gradually die off. Now, this *can* lead to a situation in which the ecosystem drives itself extinct: the foxes eliminate all the rabbits, and then they all starve and die. This certainly can happen in the wild. But suppose the foxes don’t quite manage to kill off all the rabbits—some of them elude their predators. When there is a big fox population but few rabbits, there is not enough

---

*The same is true of rabbits, of course, but because rabbits multiply so fast, we do not need to incorporate this into the picture.

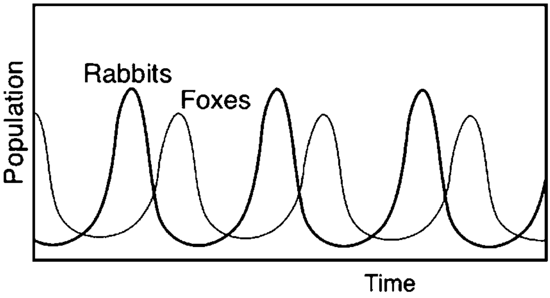

**Fig. 3.3:** Alfred Lotka's mathematical treatment of predator–prey populations gives rise to population sizes that oscillate out of step with one another.

food for all the foxes and so their population crashes. This gives the few remaining rabbits some respite, and their numbers begin to rise again. We are back to the situation we started with: lots of rabbits and few foxes. Then step 2 kicks in again and the fox population rises while the rabbits decline. Within a certain range of value of the relative rates of the three steps in the process, the system undergoes regular oscillations in the numbers of both foxes and rabbits, with the fox population peaking almost perfectly out of step with the rabbits (Fig. 3.3).

This is precisely the scheme that Lotka used to show how sustained chemical oscillations might occur. He replaced rabbits and foxes with chemical compounds; for example, suppose that a reaction of compound G[rass] with compound R[abbits] leads to more of R (an autocatalytic step). We can write this as

1. G + R → more R

In the second step of the process, R reacts with compound F[oxes] to make more F (another autocatalytic step):

2. R + F → more F

Finally, compound F decays spontaneously into another compound, D[ead foxes]:

3. F → D

This is the same scheme as that for the case of rabbits and foxes, and will likewise produce oscillations in the concentrations of reagents

R and F. Now suppose that R is coloured red and F is blue—then we have a process that changes colour back and forth, just like the BZ reaction.

The crucial point to appreciate is that the oscillations are sustained by a steady influx of materials or energy. In the case of the rabbits and foxes, the rabbits can go on multiplying only so long as there is grass. This keeps growing because of the influx of solar energy, water, and so forth. And the system never gets clogged up with the carcasses of dead foxes, because these decay and return to the soil. The ecosystem is like a CSTR in which materials (compound G) are constantly added and waste (compound D) is constantly removed. This is what prevents the system from reaching a static, unchanging equilibrium. Instead, it reaches a dynamic *steady state*—there is no stasis, but the same thing keeps repeating.

The BZ reaction is not exactly like Lotka's scheme: it is a great deal more complex, involving at least 30 different chemical compounds and many steps. When Western scientists learned about the reaction at an international conference in Prague in 1968 that brought them together with the Soviets, the chemists Richard Field, Endre Körös, and Richard Noyes at the University of Oregon set about deducing its mechanism. By 1972 they had devised a somewhat simplified scheme which accounted for the oscillations. Two years later, Field and Noyes pared this model down to an even simpler one, called the Oregonator (the 'Oregon oscillator'), in which there are just five steps involving six chemical compounds, most of them containing bromine and oxygen. Just one of these steps is autocatalytic, and rather oddly, none of them involves malonic acid (recall that the 'end result' of the BZ reaction is to add bromine atoms to this organic acid). The bromination of malonic acid does not feature in the oscillatory part of the process—it is simply a side reaction induced by the 'output' of the Oregonator, a chemical species called hypobromite.

The key feature of the Oregonator is that it has two branches—two distinct sets of reactions—one of which induces the red colour and one the blue. The system switches back and forth between these branches as they rise to dominance and then exhaust themselves, rather as Lotka's ecosystem switches between a dominance of rabbits and of foxes.

The simplest way to record the BZ oscillations is to merely keep note of the colour changes: red–blue–red–blue and so on. A more precise way is to measure the concentrations of the various compounds that

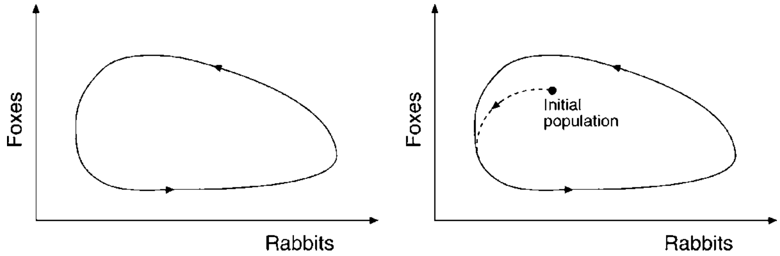

**Fig. 3.4:** Oscillations of two interacting populations can be depicted as limit cycles (*a*). If the population sizes start from a position off the limit cycle, they are soon drawn onto it (*b*). In this way, the limit cycle is said to act as an attractor.

wax and wane, like the numbers of rabbits and foxes (as in Fig. 3.3). But there is yet another way, too, which brings out a more profound feature of what we might call the 'mathematical form' of this process. Rather than plotting a graph of the concentration of a chemical species over time, we could plot how its concentration changes in relation to that of one of the other oscillating species. It is clear that the fox population is high when the rabbit population is low and vice versa, but in fact there are also times when both are medium-sized and either rising or falling. If we plotted the number of foxes against the number of rabbits in Fig. 3.3, we would get something like Fig. 3.4*a*. Each cycle of oscillation corresponds to a single circuit of this loop, which is called a limit cycle. We would get a similar limit cycle by plotting the concentration of two of the oscillating compounds in the BZ reaction. If we start the reaction off with concentrations of these reagents that lie off the limit cycle, the concentrations will evolve along a trajectory that takes them onto the loop (Fig. 3.4*b*). For this reason, the limit cycle is said to be an *attractor* of the system.

## Going places

The oscillating BZ reaction can thus be considered to display a periodic pattern of sorts, but it is a pattern in *time* (which is why the reaction has been dubbed a 'time crystal'). The oscillatory process can, however, also generate patterns in space. So far, I have talked only about a well stirred medium, which has a uniform composition throughout the reaction vessel at any instant. But if the reaction is carried out without stirring, small variations in the concentrations of the reagents are likely to arise from place to place. This is true of any chemical reaction, and normally does not lead to anything remarkable.

For an autocatalytic process like the BZ reaction, however, small variations can make a big difference. The positive feedback can have the effect of blowing up small differences into big ones. In particular, it means that one region of the BZ mixture can get flipped onto a different branch from that of the surrounding regions. The blue branch can appear in a sea of red. Then we have a mixture in which the colour changes from place to place.

What took chemists by surprise is that these colour variations do not take the form of a random patchwork of red and blue. Instead, we find complex, orderly, and rather beautiful patterns emerging. In a shallow dish of the BZ mixture, concentric rings or twisting spirals radiate outwards from a central source like ripples (Fig. 3.5 and Plate 2). The chemical oscillations give rise to moving *chemical waves*.

These patterns were first described by the German scientist Heinrich Busse in 1969, although it was not until the following year that Zhabotinsky and Zaikin correctly identified them as chemical waves—travelling 'fronts' of chemical change. It is not hard to see

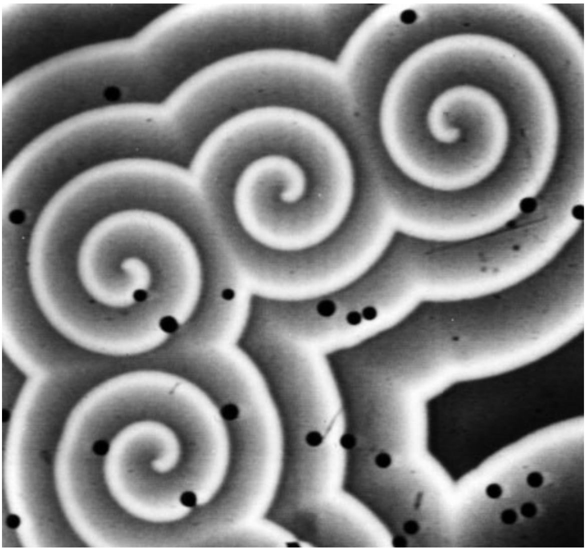

**Fig. 3.5:** Chemical waves in the Belousov–Zhabotinsky reaction. Here I show examples of spiral waves, but in general both spirals and concentric target patterns may appear (see Plate 2). (Photo: Stefan Müller, University of Magdeburg.)

how a fluctuation in the relative concentrations of the reacting species might shift the BZ reaction from one branch to the other; but why does this disturbance then radiate outwards as an organized wave?

Imagine that a small region of the solution has been tipped from the red to the blue branch. Because of the autocatalysis, this blue region expands from its origin as the molecules diffuse into the 'red' solution and induce tipping there. The rate of diffusion is equal in all directions, and so this expanding wavefront is circular. The notion that autocatalytic reactions can generate propagating chemical waves in fact predates any recognition of oscillating reactions. In 1900 the German physical chemist Wilhelm Ostwald showed that, when he pricked the dark surface of oxidized iron immersed in acid with a zinc needle, this triggered an electrochemical process that changed the colour of the surface coating, which propagated away from the point of contact at high speeds. From the 1920s some researchers studied this system as a simple analogue of nerve impulses, which are also electrochemical waves. Meanwhile, in 1906 Robert Luther, the director of the Physical Chemistry Laboratory in Leipzig, announced the discovery of chemical waves in autocatalytic reactions to an audience of German chemists in Dresden. Some were sceptical, until Luther demonstrated the process before their eyes, projecting an image of a chemical wave on to a screen.

Luther pointed out that the waves depend on competition between the autocatalytic process and the diffusion of molecules. Autocatalysis can rapidly exhaust the available resources: if the grass doesn't grow fast enough in the first step of Lotka's scheme, the rabbit population might crash through lack of food even without the foxes preying on them. Luther said that the same can happen in the corresponding chemical process—except that here the 'grass', which we called reagent G, can diffuse into the growing rabbit colony from the surroundings to replenish the supply. Because of this delicate balance between the rate of reaction and the rate of diffusion, chemical waves are said to be a *reaction–diffusion* phenomenon. After Luther's pioneering studies, the theory of reaction–diffusion systems was placed on a firm mathematical footing in the 1930s by the Russian mathematician Andrei Kolmogoroff and the English geneticist Ronald Fisher. Like Lotka and Volterra, Fisher was interested in these systems because of their relevance to population dynamics: he was investigating the rate at which an advantageous gene would spread through a population. Clearly, biologists shared none of the hesitation that chemists displayed in assimilating these ideas about complex dynamics and pattern formation.*

Autocatalysis and the reaction–diffusion process account for how an expanding reaction front can become a *wave*, where the concentration of a chemical species rises and then falls again as the wave passes. Within the wavefront, the autocatalytic process takes hold and then rapidly exhausts itself. All the while, the region just in front of the wave remains 'ripe' for colonization by the diffusing autocatalytic species—and so the wave progresses outwards. But in the BZ medium this happens not once but repeatedly: once the source of the first wave is established, it goes on discharging subsequent waves at regular intervals, giving rise to a well defined wavelength (the distance between successive wavefronts).

This, as you have probably guessed already, is because the BZ reaction is not simply autocatalytic but oscillatory. The wave source region is like a tiny flask of BZ mixture, blinking blue and red—but this 'flask' has no boundaries, and so the transformation propagates. You might wonder why all of the mixture does not behave the same way, so that it becomes a chaos of wave sources. The reason is that, once the first wavefront has passed through a part of the mixture, that region is 'enslaved' to the 'pacemaker' at the wave's origin. Behind the wavefront, the medium has been 'exhausted' by the passing wave and cannot switch branches again until the oscillatory cycle has played itself out—and it is at precisely this point that the next wave arrives. Each wave source expands its territory until it encounters a wavefront from another source. When the waves collide, they annihilate one another. This is because one wave cannot excite the region behind the other wavefront, which is in the 'exhausted' phase of the cycle. So the collisions of waves create fixed, stationary boundaries between territories commanded by each pacemaker.

---

*This might seem strange, given the aversion to mathematics and abstract theory that biologists have often shown, from D'Arcy Thompson's time to our own. But ecologists, population biologists and to some extent neuroscientists have always been unusual in this regard, embracing mathematics in a way that biologists working at the level of cells and molecules have not. Lotka's 1924 book *Elements of Physical Biology* was the first exposition on what later became known as mathematical biology (its 1956 reprint used this term instead), while Fisher made important contributions to pure statistics as well as population genetics. We will encounter several other examples in later chapters, and in the other books in this series, of biologists of this persuasion blazing a trail into the world of pattern formation well in advance of physical scientists.

An unstirred BZ mixture can thus be considered to consist of three types of region. One is at the wavefront itself, where autocatalysis induces branch switching and a change of colour. Here the medium is said to be ‘excited’, rather like the electrical surge of a nerve impulse. In front of the wave, the medium is ripe for excitation: it is in a ‘receptive’ state. And behind the wavefront, the medium is exhausted or ‘refractory’, resistant to further excitation until the cycle has run its course.

A medium that has the potential to adopt these three states is called ‘excitable’, and is liable to experience the circular, periodic travelling waves characteristic of the BZ mixture.* The ingredients of an excitable medium are rather basic and generic—they say nothing about the particular chemical reagents involved. In fact, they do not even specify a *chemical* process at all. Researchers have studied computer models of a sort of general-purpose excitable medium which is represented as a checkerboard lattice of little compartments or cells, each of which interacts with those around it. The ‘rules’ of the model are:

1. Each cell can be in either an excited, a receptive, or a refractory state.
2. Excited cells become refractory after a certain length of time, and must stay that way for a fixed period until returning to the receptive state.
3. Receptive cells are transformed to the excited state if a certain proportion of their neighbours are excited.

This sort of model, in which an array of cells adopt specific states contingent on the states of their neighbours, is called a *cellular automaton*, reflecting the fact that the cells’ behaviour is conditioned by an automatic, knee-jerk response to those around it. It is an extraordinarily versatile way to model systems that consist of many interacting components, and I shall draw on cellular-automaton schemes repeatedly in the following pages. The cellular automaton for the BZ reaction captures the essential characteristics of an excitable medium—but there is nothing in the prescription that gives any hint of the kinds of patterns that might arise. Yet when this model is run on a computer, it produces just the kind of target and spiral patterns seen in the real BZ medium (Fig. 3.6).

---

*In experiments on the BZ system, the reagents are generally infused in a layer of gel. This slows down the rates of diffusion, and makes the chemical waves more stable and less sensitive to disturbances, so that they have a smoother, regular shape.

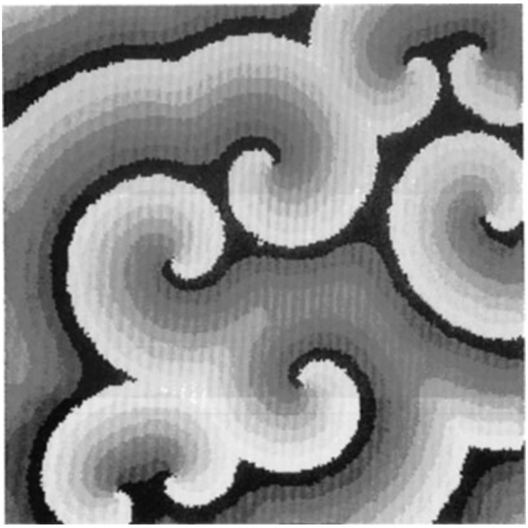

**Fig. 3.6:** The wave patterns of the BZ reaction are mimicked in a mathematical model called a cellular automaton that takes no account of any chemical details, but simply represents the mixture as a grid of cells that can be 'excited' by receiving stimuli from their neighbours. (Image: Mario Markus and Benno Hess, Max Planck Institute for Molecular Physiology, Dortmund.)

This shows that the wave pattern has nothing to do with any features of bromate or malonic acid or any of the other ingredients. It is to be expected for any system that has the characteristics of an excitable medium. The patterns are *universal*.

The account I have given so far explains why we should expect concentric target waves, but it is not clear why spiral waves are generated as well. The spirals are actually 'mutant targets', created by a disruption of the circular wavefront. Such perturbations can happen by accident, for example if there is some impurity such as a dust particle in the reaction medium; or they can be induced on purpose, for instance by blowing air on to the wavefront through a narrow pipe. At the break in the circular wave, the ends curl up and become sources of spirals.

If the BZ mixture is infused not into a thin layer of gel but into a thick slab, then the chemical waves may propagate in three dimensions. The patterns are then more complex. A spiral wave, for example, becomes a three-dimensional form called a scroll wave (Fig. 3.7). Cross-sections of scroll waves look like concentric target waves in one plane, and like pairs of counter-rotating spiral waves in another. These patterns were first seen in BZ mixtures in the 1970s.