3

PHOTOSYNTHESIS

Summoned by the Sun

Imagine a world without photosynthesis. It wouldn't be green, for a start. Our emerald planet reflects the glory of plants and algae, and ultimately their green pigments, which absorb light for photosynthesis. First among pigments is the marvellous transducer that is chlorophyll, which steals a beam of light and conjures it into a quantum of chemical energy, driving the lives of both plants and animals.

The world probably wouldn't be blue either, for the azures of the heavens and the marines of the oceans depend on clear skies and waters, cleansed of their haze and dust by the scouring power of oxygen. And without photosynthesis there would be no free oxygen.

In fact there might not be any oceans either. Without oxygen there is no ozone; and without that, there is little to cut down the searing intensity of ultraviolet rays. These split water into oxygen and hydrogen. The oxygen is formed slowly and never builds up in the air; instead it reacts with iron in the rocks, turning them a rusty-red colour. And hydrogen, the lightest of gases, evades the tug of gravity and slips away into space. The process may be slow but it is also inexorable: the oceans bleed into space. Ultraviolet radiation cost Venus its oceans, and maybe Mars too.

So we don't need much imagination to picture a world without photosynthesis: it would look a lot like Mars, a red dusty place, without oceans, and without any overt signs of life. Of course, there is life without

photosynthesis, and many astrobiologists seek it on Mars. But even if a few bacteria are found hiding beneath the surface, or buried in an icecap, the planet itself is dead. It is in near-perfect equilibrium, a sure sign of inertia. It could never be mistaken for Gaia.

Oxygen is the key to planetary life. No more than a waste product of photosynthesis, oxygen really is the molecule that makes a world. It is let loose by photosynthesis so fast that it finally overwhelms the capacity of a planet to swallow it up. In the end, all the dust and all the iron in the rocks, all the sulphur in the seas and methane in the air, anything that can be oxidised is oxidised, and free oxygen pours into the air and the oceans. Once there, oxygen puts a stop to the loss of water from the planet. Hydrogen, when released from water, inevitably bumps into more oxygen before it finds its way out into space. Swiftly it reacts to form water again, which now rains back down from the heavens, drawing to a halt the loss of the oceans. And when oxygen accumulates in the air, an ozone shield forms, ablating the searing intensity of the ultraviolet rays, and making the world a more habitable place.

Oxygen doesn't just rescue a planet's life: it energises all life, and makes it big. Bacteria can do perfectly well without oxygen: they have an unparalleled skill at electrochemistry, they are able to react together virtually all molecules to glean a little energy. But the sum total of energy that can be derived from fermentation, or by reacting two molecules like methane and sulphate together, is negligible in comparison with the power of oxygen respiration – literally the burning up of food with oxygen, oxidising it fully to carbon dioxide and water vapour. Nothing else can provide the energy needed to fuel the demands of multicellular life. All animals, all plants, all of them depend on oxygen for at least part of their life cycle. The only exception that I'm aware of is a microscopic (but multicellular) nematode worm that somehow gets along in the stagnant oxygen-free depths of the Black Sea. So a world without free oxygen is microscopic, at least at the level of individual organisms.

Oxygen contributes to large size in other ways too. Think of a food chain. The top predators eat smaller animals, which might in turn eat insects, which eat smaller insects, which live on fungus or leaves. Five or six levels in a food web are not uncommon. At each step energy is wasted, for no form of respiration is ever 100 per cent efficient. In fact, oxygen respiration is about 40 per cent efficient, while most other forms of respiration (using iron or sulphur instead of oxygen, for example) are less than 10 per cent efficient. This means that, without using oxygen, the energy available dwindles to 1 per cent of the initial input in only two levels, whereas with oxygen it takes six levels to arrive at the same point. That in turn means that long food chains are only feasible with oxygen respiration. The economy of the food chain means that predators can operate in an oxygenated world, but predation as a lifestyle just doesn't pay without oxygen.

Predation escalates size, of course, driving arms races between predator and prey. Shells combat teeth, camouflage tricks the eye; and size intimidates both hunter and hunted. With oxygen, then, predation pays; and with predators size pays. So oxygen makes large organisms not just feasible but also probable.

It also helps build them. The protein that gives animals their tensile strength is collagen. This is the main protein of all connective tissues, whether calcified in bones, teeth and shells, or 'naked' in ligaments, tendons, cartilage and skin. Collagen is by far the most abundant protein in mammals, making up a remarkable 25 per cent of total body protein. Outside the vertebrates, it is also the critical component of shells, cuticles, carapaces and fibrous tissues of all sorts – the 'tape and glue' of the whole animal world. Collagen is composed of some unusual building blocks, which require free oxygen to form cross-links between adjacent protein fibres, giving the overall structure a high tensile strength. The requirement for free oxygen means that large animals, protected with shells or strong skeletons, could only evolve when atmospheric oxygen levels were high enough to support collagen production – a factor that might have contributed to the abrupt appearance of large animals in the fossil record at the beginning of the Cambrian period, some 550 million years ago, soon after a big global rise in atmospheric oxygen.

The need for oxygen to make collagen may seem no more than an accident; if not collagen why not something else with no requirement for free oxygen? Is oxygen necessary to give strength or just a random ingredient that happened to be incorporated and then forever remained part of the recipe? We don't really know, but it's striking that higher plants, too, need free oxygen to form their structural support, in the shape of the immensely strong polymer lignin, which gives wood its flexible strength. Lignin is formed in a chemically haphazard way, using free oxygen to form strong cross-links between chains. These are very difficult to break down, which is why wood is so strong and why it takes so long to rot. Eliminate lignin from trees – a trick that manufacturers of paper have tried, as they need to remove it laboriously from wood pulp to make paper – and the trees slump to the ground, unable to sustain their own weight even in the lightest breeze.

So without oxygen there would be no large animals or plants, no predation, no blue sky, perhaps no oceans, probably nothing but dust and bacteria. Oxygen is without a doubt the most precious waste imaginable. Yet not only is it a waste product, it is also an unlikely one. It is quite feasible that photosynthesis could have evolved here on earth, or Mars, or anywhere else in the universe, without ever producing any free oxygen at all. That would almost certainly consign any life to a bacterial level of complexity, leaving us alone as sentient beings in a universe of bacteria.

✦

One reason why oxygen might never have accumulated in the air is respiration. Photosynthesis and respiration are equal and opposite processes. In a nutshell, photosynthesis makes organic molecules from two simple molecules, carbon dioxide and water, using sunlight to provide the energy needed. Respiration does exactly the opposite. When we burn organic molecules (food) we release carbon dioxide and water back into the air; and the energy released is what powers our lives. All our energy is a beam of sunlight set free from its captive state in food.

Photosynthesis and respiration oppose each other not just in the details of their chemistry, but also in global accounting. If there was no respiration – no animals, fungi and bacteria burning up plant food – then all the carbon dioxide would have been sucked out of the atmosphere long ago, converted into biomass. Everything would then more or less grind to a halt, bar the trickle of carbon dioxide set free by slow decay or volcanoes. But this is far from what really happens. What really happens is that respiration burns all the organic molecules put away by plants: on a geological timeframe, plants disappear in a puff of smoke. This has one profound consequence. All the oxygen put in the air by photosynthesis is taken out again by respiration. There is a long-term, unchanging, never-ending equilibrium, the kiss of death for any planet. The only way that a planet can gain an oxygen atmosphere – the only way it can escape the dusty red fate of Mars – is if a little plant matter is preserved intact, immune to the elements and to life's ingenuity in finding ways of breaking it down for energy. It must be buried.

And so it is. Preserved plant matter is buried as coal, oil, natural gas, soot, charcoal or dust, in rocks deep in the bowels of the earth. According to the ground-breaking geochemist Robert Berner, recently retired from Yale, there is around 26,000 times more 'dead' organic carbon trapped in the earth's crust than in the entire living biosphere. Each atom of carbon is the antithesis of a molecule of oxygen in the air. For every atom of carbon that we dig up and burn as fossil fuel, a molecule of oxygen is stripped out of the air, and converted back to carbon dioxide, with serious, albeit unpredictable, consequences for climate. Luckily we will never deplete the world's oxygen supply by burning fossil fuels, even if we do play havoc with the climate: the vast majority of organic carbon is buried as microscopic detritus in rocks like shales, inaccessible to human industry, or at least economic industry. So far, despite our vainglorious efforts to burn all the known reserves of fossil fuels, we have lowered the oxygen content of the air by a mere two or three parts per million, or about 0.001 per cent.[1]

But this vast reservoir of buried organic carbon is not formed continuously – it has been buried in fits and starts over the geological aeons. The norm is very close to an exact balance, in which respiration cancels out photosynthesis (and erosion cancels out any burial), so there is next to no net burial. This is why oxygen levels have remained at around 21 per cent for tens of millions of years. In deep geological time, though, on rare occasions, things were very different. Perhaps the most striking example is the Carboniferous period, 300 million years ago, when dragonflies as big as seagulls flapped through the air and millipedes a metre long crawled the undergrowth. These giants owed their very existence to the exceptional rate of carbon burial in Carboniferous times, whose huge coal reserves give the era its name. As carbon was buried

beneath the coal swamps, oxygen levels soared above 30 per cent, giving some creatures an opportunity to grow far beyond their normal bounds of size – specifically, animals that rely on the passive diffusion of gases down tubes or across the skin, like dragonflies, rather than the active ventilation of lungs.[2]

What was behind the unprecedented rate of carbon burial in Carboniferous times? A variety of accidental factors, almost certainly. The alignment of the continents, the wet climate, the great flood plains; and perhaps most importantly the evolution of lignin, which gave rise to large trees and sturdy plants capable of colonising large areas of the landmass. Lignin, tough to break down for fungi and bacteria even today, seems to have been an insurmountable challenge soon after its evolution. Rather than being broken down for energy, it was buried intact on a vast scale, and its antithesis, oxygen, flooded the air.

Geological accidents colluded on two other occasions to force up oxygen levels, both perhaps the outcome of global glaciations called 'snowball earths'. The first great rise in oxygen levels, around 2,200 million years ago, followed hard on the heels of a period of geological upheavals and global glaciation around that time; and a second period of global glaciations, around 800 to 600 million years ago, also seems to have pushed up oxygen levels. Such calamitous global events probably altered the balance of photosynthesis to respiration, and of burial to erosion. As the great glaciers melted and the rains fell, minerals and nutrients (iron, nitrates and phosphates) that had been scoured from the rocks by ice were washed into the oceans, causing a great bloom of photosynthetic algae and bacteria, similar to, but far greater than, those caused today by fertilisers. Not only would such a run-off induce a bloom, it would also tend to bury it: the dust, dirty ice and grit washed into the oceans mixed with blooming bacteria and settled out, burying carbon on an unprecedented scale. And with it came a lasting global rise in oxygen.

So there is a sense of the accidental about the oxygenation of our planet. This sense is reinforced by the absence of change for long periods otherwise. From 2,000 million to around 1,000 million years ago – a period geologists call the 'boring billion' – almost nothing of note seems to have happened. Oxygen levels remained steady and low throughout this period, as indeed they did at other times for hundreds of millions of years. Stasis is the default, while episodes of geological restlessness wreak lasting change. Such geological factors might intervene on other planets too; but tectonic movements and active volcanism seem to be necessary to bring about the accidental conjugations needed for oxygen to accumulate. It is not beyond the bounds of possibility that photosynthesis evolved long ago on Mars, but that this small planet, with its shrinking volcanic core, could not sustain the geological flux required for oxygen to accumulate, and later expired on a planetary scale.

✦

But there is a second and more important reason why photosynthesis need not lead to an oxygen atmosphere on a planet. Photosynthesis itself may never turn upon water as a raw material. We are all familiar with the form of photosynthesis that we see around us. Grasses, trees, seaweeds, all operate in fundamentally the same way to release oxygen – a process known as 'oxygenic' photosynthesis. But if we take several steps back and consider bacteria, there are many other options. Some relatively primitive bacteria make use of dissolved iron or hydrogen sulphide instead of water. If these sound like implausible raw materials to us, it is only because we have become so inured to our oxygenated world – the product of 'oxygenic' photosynthesis – that we struggle to imagine conditions on the early earth when photosynthesis first evolved.

We also struggle to grasp the counterintuitive, but in fact simple, mechanism of photosynthesis. Let me give an example, which I suspect, perhaps unfairly, illustrates the general perception of photosynthesis. This is Primo Levi, from his lovely book *The Periodic Table*, published in 1975 and voted the 'best popular science book ever' by an audience (including me) at the Royal Institution in London, in 2006:

> Our atom of carbon enters the leaf, colliding with other innumerable (but here useless) molecules of nitrogen and oxygen. It adheres to a large and complicated molecule that activates it, and simultaneously receives the decisive message from the sky, in the flashing form of a packet of solar light: in an instant, like an insect caught by a spider, it is separated from its oxygen, combined with hydrogen and (one thinks) phosphorus, and finally inserted in a chain, whether long or short does not matter, but it is the chain of life.

Spot the mistake? There are actually two, and Levi ought to have known better, for the true chemistry of photosynthesis had been elucidated forty years earlier. A flashing packet of solar light does not activate carbon dioxide: it can be activated just as well in the middle of the night, and indeed is never activated by light, even in the brightest sunshine. Nor is carbon separated in an instant from its oxygen. Oxygen remains stubbornly bound to its carbon. Underpinning Levi's account is the common, but plain wrong, assumption that the oxygen released by photosynthesis comes from carbon dioxide. It does not. It comes from water. And that makes all the difference in the world. It is the first step to understanding how photosynthesis evolved. It is also the first step to solving the energy and climate crises of our planet.

The packets of solar energy used in photosynthesis split water into hydrogen and oxygen: the same reaction that occurs on a planetary scale when the oceans bleed into space, driven away by the blast of ultraviolet radiation. What photosynthesis achieves – and what we have so far failed to achieve – is to come up with a catalyst that can strip the hydrogen from water with a minimal input of energy, using gentle sunlight rather than searing ultraviolet or cosmic rays. So far, all our human ingenuity ends up consuming more energy in splitting water than is gained by the split. When we succeed in mimicking photosynthesis, with a simple catalyst that gently prises hydrogen atoms from water, then we will have solved the world's energy crisis. Burning that hydrogen would comfortably supply all the world's energy needs, and regenerate water as the only waste: no pollution, no carbon footprint, no global warming. Yet this is no easy task, for water is a marvellously stable combination of atoms, as the oceans attest; even the most furious storms, battering cliffs, don't break water into its component atoms. Water is at once the most ubiquitous and unattainable raw material on our planet. The modern mariner might muse on how to power his boat with water and a splash of sunshine. He should ask the green scum floating on the waves.

The same problem, of course, faced the remote ancestors of that scum, the ancestors of today's cyanobacteria, the only form of life on our planet to have chanced upon the trick of splitting water. The strange thing is that cyanobacteria split water for exactly the same reason that their bacterial relatives split hydrogen sulphide or oxidise iron: they want the electrons. And on the face of it, water is the last place to find them.

Photosynthesis is conceptually simple: it's all about electrons. Add a few electrons to carbon dioxide, along with a few protons to balance out charges, and, hey presto, there you have it – a sugar. Sugars are organic molecules: they are Primo Levi's chain of life and the ultimate source of all our food. But where do the electrons come from? With a little energy from the sun they can come from more or less anywhere. In the case of the familiar 'oxygenic' form of photosynthesis, they come from water; but in fact it's far easier to strip them from other compounds less stable than water. Take electrons from hydrogen sulphide and instead of releasing oxygen into the air you deposit elemental sulphur – biblical brimstone. Take them from iron dissolved in the oceans (as ferrous iron) and you get rusty-red ferric iron, which settles out as new rocks – a process that might once have been responsible for the vast 'banded-iron formations' found around the world, and today the largest remaining reserves of low-grade iron ore.

These forms of photosynthesis are marginal in today's oxygen-rich world, simply because the raw materials, hydrogen sulphide or dissolved iron, are rarely found in sunny well-aerated waters. But when the earth was young, before the rise of free oxygen, they would have been by far the easiest source of electrons, and they saturated the oceans. This raises a paradox, whose resolution is fundamental to understanding how photosynthesis first evolved. Why switch from a plentiful and comfortable source of electrons to something far more problematic, water, whose waste product, oxygen, was a toxic gas capable of causing grievous bodily harm to any bacteria that produced it? The fact that, given the power of the sun and a clever catalyst, water is far more abundant than either, is beside the point, for evolution has no foresight. So too is the fact that oxygenic photosynthesis transformed the world; the world cares not a whit. So what kind of environmental pressure, or mutations, could have driven such a shift?

The facile answer, which you'll find in plenty of textbooks, is that the raw

materials ran out: life turned to water because there were no easy alternatives left, just as we might turn to water when we run out of fossil fuels. But this answer can't be true: the geological record makes it plain that 'oxygenic' photosynthesis evolved long before – more than a billion years before – all these raw materials ran out. Life was not forced into a corner.

Another answer, only now emerging, lies hidden in the machinery of photosynthesis itself, and is altogether more beautiful. It is an answer that combines chance and necessity, an answer that shines the light of simplicity on one of the most convoluted and complicated extractions in the world.

✦

In plants the business of electron extraction takes place in the chloroplasts. These are the minute green structures found in the cells of all leaves, all blades of grass, imparting their own green to the leaves as a whole. The chloroplasts are named from the pigment that gives them their colour in turn. This is chlorophyll, which is responsible for absorbing the energy of the sun in photosynthesis. Chlorophyll is embedded in an extraordinary membrane system that makes up the inside of the chloroplasts. Great stacks of flattened disks, looking for all the world like an alien power station in a science fiction movie, are connected to each other via flying tubes, which criss-cross the vertiginous spaces at all angles and heights. In the disks themselves, the great work of photosynthesis takes place: the extraction of electrons from water.

If the extraction of electrons from water is difficult, plants make an extraordinary meal of it. The complexes of proteins and pigments are so vast, in molecular terms, that they amount to a small city. Altogether they form into two great complexes, known as Photosystem I and Photosystem II, and each chloroplast contains thousands of such photosystems. Their job is to catch a beam of light, and transform it into living matter. Working out how they do it has taken the best part of a century, and took some of the most elegant and ingenious experiments ever done. This, sadly, is not the place to discuss them.[3] Here we need concern ourselves only with what we have learned, and what that has to say about the invention of photosynthesis.

The conceptual heart of photosynthesis, the roadmap that makes sense of it all, is known as the 'Z scheme', a formulation that fascinates and horrifies biochemistry students in equal measure. First laid out by the brilliant but diffident Englishman Robin Hill in 1960, the Z scheme describes the 'energy profile' of photosynthesis. Hill's utterances were notoriously gnomic. Wishing not to cause offence by labouring the obvious, even those in his own lab were surprised when his hypothesis appeared in *Nature* in 1960, having had little idea what he was working on. In fact the Z scheme was not based on Hill's own work, which had been of primary importance, but rather was distilled from a number of puzzling experimental observations. Foremost among these was a curious matter of thermodynamics. Photosynthesis, it turned out, produced not just new organic matter but also ATP, the 'energy currency' of life. Quite unexpectedly, these two seemed always to be coupled: the more organic matter produced by photosynthesis, the more ATP, and vice versa (if the amount of organic matter falls, so too does ATP production). The sun apparently provides two free lunches in synchrony. Robin Hill, remarkably, had the insight to grasp the whole mechanism of photosynthesis from this single fact. Genius, it is said, is the ability to see the obvious before anyone else.[4]

And yet – typically for anything associated with Hill – even the term 'Z scheme' is gnomically misleading. The Z should really be rotated through 90 degrees to become an 'N'; then it would reflect more accurately the energy profile of photosynthesis. Picture the first upstroke of the 'N' as a vertical uphill reaction: energy must be supplied to make it work. The diagonal downstroke of the 'N' is then a downhill reaction – it releases energy that can be captured and stored in the form of ATP. The final upstroke is again an uphill reaction, which requires an input of energy.

In photosynthesis, the two photosystems – Photosystem I and II – lie at the two bottom points of the 'N'. A photon of light hits the first photosystem and blasts an electron up to a higher energy level; the energy of this electron then cascades down in a series of small molecular steps, which provide the energy needed to make ATP. Back at a low energy level, the electron arrives at the second photosystem, where a second photon blasts it up a second time to a higher energy level. From this second high point, the electron is ultimately transferred to carbon dioxide, in the first step of making a sugar. One helpful

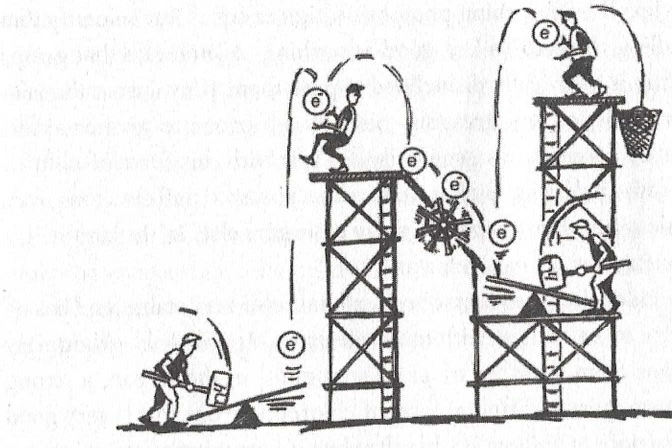

Figure 3.1 Cartoon depicting the Z scheme, by Richard Walker. The energy of a photon, depicted as a mallet blow, blasts an electron to a high energy level. As the electron cascades back down to a lower energy level, some of the energy released powers work in the cell. A second photon then blasts the electron back up to an even higher energy level, where it is captured in the form of a high-energy molecule (NADPH) that later reacts with carbon dioxide to form an organic molecule.

cartoon by Richard Walker (see Fig. 3.1) depicts the process as a fairground test-of-strength game, where a punter hits a pallet with a mallet to ring a bell by forcing a metal ringer up a pole. In this case, the swing of the mallet provides the energy to blast the ringer up the pole; in the case of photosynthesis, the energy of a photon from the sun does the same job.

The Z scheme, or N scheme if you prefer, is a curiously convoluted way of going about things, but there are good technical reasons for it. It verges on the chemically impossible to couple the removal of electrons from water to the conversion of carbon dioxide into a sugar in any other way. The reason relates to the nature of electron transfer, and specifically the chemical affinity of electrons for particular compounds. Water is very stable, as we've seen: it has a high affinity for its electrons. To steal an electron from water requires huge pulling power, which is to say a very powerful oxidant. That powerful oxidan is a voracious form of chlorophyll, a molecular Mr Hyde, transformed from the meek Dr Jekyll by absorbing photons of high energy.[5] But an entity that is good at pulling tends to be less good at pushing. A molecule that grasps electrons tightly is chemically disinclined to push them away, just as the misanthropic Mr Hyde, or any grasping miser, is not prone to give away his wealth in acts of spontaneous generosity. So it is with this form of chlorophyll. When activated by light it has tremendous power to pull electrons from water, but little strength to push them away anywhere else. In the jargon, it is a powerful oxidant but a weak reductant.

Carbon dioxide raises the reverse problem. It is also very stable, and has no chemical desire to be stuffed with more electrons. It will only grudgingly accept electrons from a pusher of great strength – in the jargon, a strong reductant. This requires a different form of chlorophyll: one that is very good at pushing, and poor at pulling. Rather than being a grasping miser, it is more like a street hustler, intent on forcing dodgy goods onto vulnerable passers-by. When activated by light, this form of chlorophyll has the power to force its electrons onto another molecule that wants equally to be rid of them, its conspirator in crime and co-hustler, NADPH, and ultimately onto carbon dioxide.[6]

So there is a reason for having two photosystems in photosynthesis. No real surprise there. But the more challenging question is: how did such a complex interrelated system come to evolve? There are actually five parts to this sequence. First is the 'oxygen-evolving complex', a kind of molecular nutcracker that positions the water molecules just so to have their electrons cracked out one by one, releasing oxygen as waste. Then comes Photosystem II (rather confusingly, the two photosystems are named in reverse order, for historical reasons), which when activated by light transforms into the molecular Mr Hyde, and yanks out these electrons from the oxygen-evolving complex. Then comes an electron-transport chain, which transfers the electrons away, like rugby players passing a ball across a pitch. The electron-transport chain uses the downhill energy gradient to make a little ATP, before delivering up the same electrons to Photosystem I. Here another photon blasts them up to a high energy level again, where they are held in trust by the molecular 'hustler' NADPH, a strong pusher of electrons, which wants nothing better

than to be rid of them again. And then finally comes the molecular machinery needed to activate carbon dioxide and convert it to a sugar. Using the molecular hustler generated by Photosystem I, the conversion of carbon dioxide into a sugar is powered by chemistry rather than light, and is actually known as the dark reaction – a feature that Primo Levi failed to appreciate.

These five systems work in sequence to strip electrons from water and push them onto carbon dioxide. It's an enormously complicated way to crack a nut, but it seems to be about the only way to crack this particular nut. The great evolutionary question is how did all these complex interrelated systems come into existence, and come to be organised in exactly the right way, perhaps the only way, to make oxygenic photosynthesis work?

✦

The word 'fact' is always likely to make biologists tremble in their boots, as there are so many exceptions to every rule; but one such 'fact' is virtually certain about oxygenic photosynthesis – it only evolved once. The seat of photosynthesis, the chloroplast, is found in all photosynthetic cells of all plants and all algae. Chloroplasts are omnipresent and are obviously related to each other. They share a secret history. The clue to their past lies in their size and shape: they look like little bacteria living inside a larger host cell (see Fig. 3.2). This hint of bacterial ancestry is confirmed by the existence of independent rings of DNA in all chloroplasts. These rings of DNA are copied whenever chloroplasts divide, and passed on to the daughters in the same way as bacteria. The detailed sequence of letters in chloroplast DNA not only corroborates the link with bacteria, but also points an accusing finger at the closest living relative: cyanobacteria. Last but not least, the Z scheme of plant photosynthesis, along with all five of its component parts, is presaged exactly (if with simpler machinery) in cyanobacteria. In short, there is no doubt that chloroplasts were once free-living cyanobacteria.

Once misnamed, poetically, the 'blue-green algae', the cyanobacteria are the only known group of bacteria that can split water via the 'oxygenic' form of photosynthesis. Exactly how some of their number came to live within a larger host cell is a mystery wrapped in the shrouds of deep geological time.

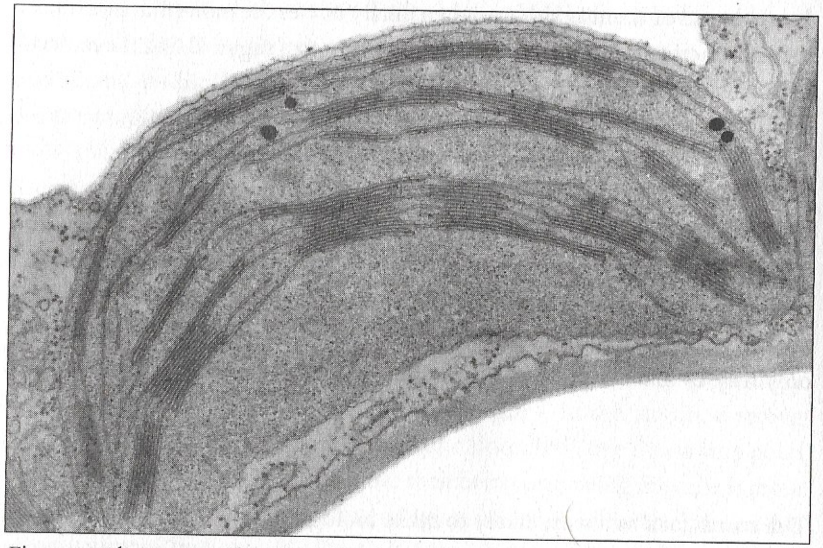

Figure 3.2 Classic view of a chloroplast from beet (*Beta vulgaris*), showing the stacks of membranes (thylakoids) where water is split apart to release oxygen in photosynthesis. The resemblance to a bacterium is not accidental: chloroplasts were once free cyanobacteria.

It undoubtedly happened more than 1,000 million years ago, but presumably they were simply engulfed one day, survived digestion (not uncommon), and ultimately proved useful to their host cell. The host, impregnated with cyanobacteria, went on to found two great empires, the algae and the plants, for all of them today are defined by their ability to live on sun and water, by way of the photosynthetic apparatus inherited from their bacterial guests.

So the quest for the origin of photosynthesis becomes the quest for the origin of cyanobacteria, the only type of bacteria to have cracked the problem of splitting water. And this is one of the most controversial, and indeed still unresolved, stories in modern biology.

Until the turn of the millennium, most researchers were persuaded, if vexed, by the remarkable findings of Bill Schopf, an energetic and combative professor of palaeobiology at the University of California, Los Angeles. From the 1980s, Schopf had discovered and analysed a number of the oldest

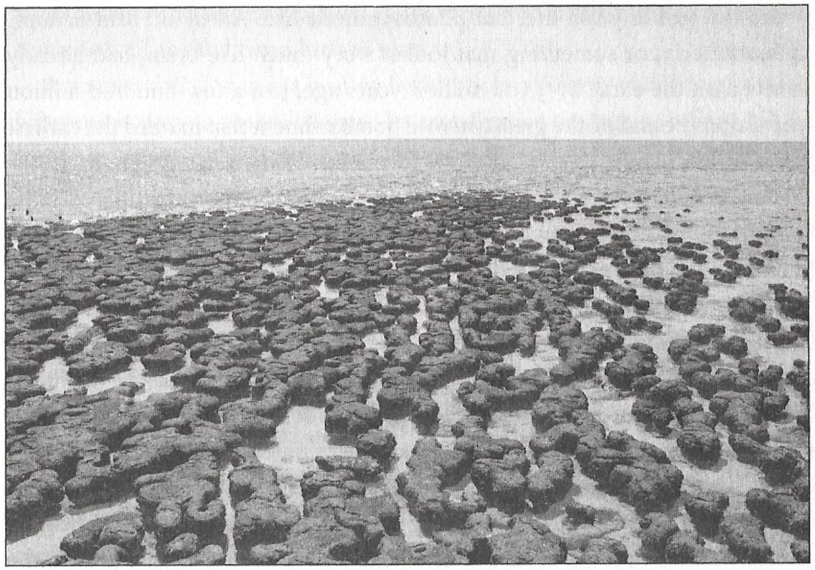

Figure 3.3 Living stromatolites in Hamelin Pool, near Shark Bay, western Australia. The pool is approximately double the salinity of the open ocean, which stifles grazers like snails and enables the cyanobacterial colonies to flourish.

fossils of life on earth – some 3,500 million years old. The word 'fossil' needs a little clarification here. What Schopf found were strings of microscopic rock capsules, which looked a lot like bacteria, and were about the right size. From their detailed structure, Schopf initially proclaimed the fossils to be cyanobacteria. These tiny microfossils were often associated with what looked like fossil stromatolites. Living stromatolites are mineralising domes that grow in incremental layers, up to a metre or so in height, which are formed by thriving communities of bacteria that encrust the buried mineral layers (see Fig. 3.3). Eventually the entire structure turns to solid rock, often strikingly beautiful in section. The outer, living layers of modern stromatolites are usually heaving with cyanobacteria, so Schopf was able to claim these ancient forms as further evidence for the early appearance of cyanobacteria. Lest there be any doubt, Schopf went on to show that these putative fossils contained remnants of organic carbon, of a sort that seemed diagnostic of life

– and not just any old life, but photosynthetic life. All in all, said Schopf, cyanobacteria, or something that looked very much like them, had already evolved on the earth by 3,500 million years ago, just a few hundred million years after the end of the great asteroid bombardment that marked the earliest years of our planet, so soon after the formation of the solar system itself.

Few people were equipped to challenge Schopf's interpretation of these ancient fossils, and the few that were also seemed convinced. Others, though, if less expert, were more sceptical. It was not easy to reconcile the early evolution of cyanobacteria – presumably belching out oxygen as a waste product, as they do today – with the first geological signs of oxygen in the atmosphere, well over a billion years later. And perhaps more seriously still, the complexity of the Z scheme made most biologists baulk at the idea that oxygenic photosynthesis could have evolved so fast. The other forms of photosynthesis, being simpler, seemed more in keeping with great antiquity. Overall, then, most people accepted that these were bacteria, perhaps photosynthetic bacteria, but there were doubts about whether they were really cyanobacteria, the pinnacle of the art.

Then Martin Brasier, professor of palaeobiology at Oxford, stepped into the ring, in what turned into one of the great fights of modern paleontology, a science in any case noted for the passion of its protagonists, and the elasticity of much of its evidence. Most researchers interested in the early fossils relied on specimens deposited in the Natural History Museum, London, but Brasier returned to the geological setting where Schopf had originally dug up his fossils, and expressed shock. Far from being the shallow, tranquil seafloor posited by Schopf, the entire region was shot through with geothermal veins, evidence of a tumultuous geological past, said Brasier. Schopf had handpicked his specimens to make his case, Brasier went on, and had concealed other specimens, superficially similar, but patently not biological; they were probably all formed by the action of scalding water on mineral sediments. The stromatolites too, he said, were formed by geological processes, not bacteria, and were no more mysterious than ripples in the sand. And the organic carbon had no microscopic structure at all, making it quite indistinguishable from the inorganic graphite found in many geothermal settings. Finally, as if to drive a stake through the corpse of a once-great scientist, one former graduate student recalled being bullied and forced into making dubious interpretations. Schopf looked like a broken man.

But never a man to take a beating lightly, Schopf emerged fighting. Assembling more data to make his case, he met Brasier onstage at a fiery NASA spring meeting in April 2002, and the pair defended their corners. Brasier, every inch the haughty Oxford don, condemned Schopf's case as 'a truly hydrothermal performance – all heat and not much light'. Even so, the jury has not really been convinced by either side. While there is real doubt about the biological origin of the earliest microfossils, others, dating to only a hundred million years later, are less contested; and Brasier himself has put forward candidate fossils from this time. Most scientists, including Schopf, now apply more stringent criteria to verify biological provenance. The one casualty so far is the cyanobacteria, once the centrepiece of Schopf's fame. Even Schopf concedes that the microfossils are probably not cyanobacteria, or at least are no more likely to be cyanobacteria than any other type of filamentous bacteria. And so we have arrived back in the starting blocks, chastened, and with no better idea of the evolution of cyanobacteria than we had at the outset.

✦

I use this tale to illustrate just how difficult it is to fathom the depths of geological time using the fossil record alone. Even proving the existence of cyanobacteria, or at least their ancestors, does not prove that they had already stumbled upon the means of splitting water. Perhaps their ancestors relied on a more primitive form of photosynthesis. But there are other ways to mine information from deep time that may yet prove more informative. These are the secrets buried within living things themselves, both in their genes and in their physical structures, especially their protein structures.

Over the last two or three decades, the detailed molecular structures of both plant and bacterial photosystems have come under intense scrutiny, with scientists bringing to bear a great battery of techniques with daunting names, signifying no less daunting methodologies, from X-ray crystallography to electron-spin resonance spectroscopy. How these techniques work needn't concern us here; suffice to know that they have been used to map out the shapes and structures of the photosynthetic complexes in nearly, but tantalisingly not quite, atomic resolution. Even now, arguments rage at meetings, but they are arguments about details. As I write, I have recently returned from a discussion meeting at the Royal Society in London that was rich in argument about the exact location of five critical atoms in the oxygen-evolving complex. The arguments were at once fiddling and profound. Profound, because their exact position defines the strict chemical mechanism by which water is split; knowing this is the key step to solving the world's energy crisis. But fiddling, because their squabbles are about positioning these five atoms to within a space of a few diameters of an atom – a few angstroms (less than a millionth of a millimetre). To the astonishment of the older generation of researchers, there is little real disagreement about the position of all the other 46,630 atoms of Photosystem II, mapped out by Jim Barber's team at Imperial College in 2004, and more recently in even more detail.

While these few atoms have yet to be allocated their final resting places, the larger architecture of the photosystems, hinted at for more than a decade, is now plain, and speaks volumes about their evolutionary history. In 2006, a small team led by Bob Blankenship, now a distinguished professor at the University of Washington in St Louis, showed that the two photosystems are extraordinarily well conserved in bacteria.[7] Despite enormous evolutionary distances between the various groups of bacteria, the core structures of the photosystems are almost identical, to the point that they can be superimposed in space using a computer. In addition, Blankenship confirmed another link that researchers had suspected for a long time: the two photosystems (Photosystem I and II) also share core structures, and almost certainly evolved from a common ancestor, long, long ago.

In other words, there was once a single photosystem. At some point the gene became duplicated to give two identical photosystems. These slowly diverged from each other under the influence of natural selection, while retaining a close structural similarity. Ultimately, the two photosystems were yoked together in the Z scheme of cyanobacteria, and later passed on to the plants and algae in chloroplasts. But this simple narrative conceals a fascinating dilemma. Duplicating a primitive photosystem could never solve the

problem of oxygenic photosynthesis – it could never couple a strong puller to a strong pusher of electrons. Before photosynthesis could work, the two photosystems had to diverge in opposite directions, and only then could they become usefully interlinked. So the question is, what succession of events could drive them apart, only to link them together again as intimate but opposite partners, like man and woman, reunited after diverging from an egg?

The best way to find the answer is to look to the photosystems themselves. These are united in the Z scheme of cyanobacteria, but otherwise have interestingly deviant evolutionary histories. Let's put aside for the moment where the photosystems originally came from, and take a quick look at their current distribution in the bacterial world. Apart from the cyanobacteria, they are never found together in the same bacterium. Some groups of bacteria have only Photosystem I, while other groups have only Photosystem II. Each photosystem works by itself to achieve different ends; and their precise tasks give a striking insight into how oxygenic photosynthesis first evolved.

In bacteria, Photosystem I does exactly the same as it does in plants. It draws electrons from an inorganic source and forms a molecular 'street hustler' that pushes them on to carbon dioxide to make sugars. What differs is the inorganic source of electrons. Rather than water, which it can't handle at all, Photosystem I draws electrons from hydrogen sulphide or iron, both of which are far easier targets than water. Incidentally, the molecular 'hustler' formed by Photosystem I, NADPH, can also be formed by pure chemistry, for example in the hydrothermal vents we discussed in Chapter 1. Here too NADPH is used to convert carbon dioxide into sugars via a similar set of reactions. So the only real innovation of Photosystem I was to harness light to do a job that was previously done by chemistry alone.

It's also worth noting here that there is nothing particularly special about the capacity to convert light into chemistry: almost any pigment can do it. The chemical bonds in pigments are good at absorbing photons of light. When they do so an electron is zapped up to a higher energy level, and nearby molecules can easily capture it. As a result the pigment becomes photo-oxidised: it is in need of an electron to balance the books, and takes one from iron or hydrogen sulphide. This is all that chlorophyll does. Chlorophyll is a porphyrin, not dissimilar in its structure to haem, the pigment that carries oxygen in our blood. Many other porphyrins can pull off similar tricks with light, sometimes with unwelcome consequences, as in diseases like porphyria.[8] And crucially, porphyrins are among the more complex molecules that have been isolated from asteroids and synthesised in the lab under plausibly prebiotic conditions. Porphyrins, in other words, would most likely have formed spontaneously on the early earth.

In short, Photosystem I took a simple-enough pigment, a porphyrin, and coupled its spontaneous light-driven chemistry to reactions that take place in bacterial cells anyway. The outcome was a primitive form of photosynthesis that could use light to strip electrons from 'easy' sources, such as iron and hydrogen sulphide, and pass the electrons on to carbon dioxide to form sugars. Thus these bacteria use light to make food.

What about Photosystem II? The bacteria that use this photosystem use light to pull off quite a different trick. This form of photosynthesis doesn't produce organic matter. Rather, it converts light energy into chemical energy, indeed electricity, which can be used to power the cell. The mechanism is very simple. When a photon strikes a molecule of chlorophyll, one electron is zapped up to a higher energy level, as before, where it is captured by a nearby molecule. This electron is then passed hot-handed from carrier to carrier down an electron-transport chain, each time releasing a little energy, until it has returned to a low energy level. Some of the energy released in this process is captured, to make ATP. Finally, the exhausted electron is returned to the same chlorophyll that it started out from, completing the circuit. In short, light zaps an electron to a high energy level, and, as it cascades back down to a 'resting' level, the release of energy is captured as ATP, a form of energy that the cell can use. It's just a light-powered electric circuit.

How did such a circuit come to be? Again, the answer is by mixing and matching. The electron-transport chain is more or less the same as that used for respiration, which evolved in the vents as we saw in Chapter 1; it was just borrowed for an ever-so-slightly new purpose. In respiration, as we noted, electrons are stripped from food and passed, ultimately, to oxygen, to form water. The energy released is used to generate ATP. In this form of photosynthesis, exactly the same thing happens: high-energy electrons are passed

along a chain, not to oxygen, but to a 'grasping' (oxidising) form of chlorophyll. The more that the chlorophyll can 'pull' electrons (that is, the closer it is to oxygen in chemical character), the more efficient the chain will be, sucking electrons along and drawing out their energy. The great advantage is that no fuel, or food, is needed, at least to provide energy (it is needed to synthesise new organic molecules).

As a general conclusion, then, the simpler forms of photosynthesis are mosaic-like in character. Both forms plugged a new transducer, chlorophyll, into existing molecular machinery. In one case, this machinery converts carbon dioxide into sugars; in the other, it produces ATP. As to chlorophyll, similar porphyrin pigments probably formed spontaneously on the early earth, and natural selection did the rest. In each case, small changes in the structure of chlorophyll alter the wavelength of light absorbed, and so the chemical properties. All these changes alter the efficiency of processes that happen spontaneously, albeit initially much more wastefully. The natural outcome is to produce a 'grasping miser' form of chlorophyll for ATP synthesis in some footloose bacteria and a 'street hustler' type of chlorophyll to make sugars in bacteria living close to supplies of hydrogen sulphide or iron. But we're still left with the bigger question: how did it all come to be tied together in the Z scheme of cyanobacteria, to split the ultimate fuel, water?

The short answer is we don't know for sure. There are ways of finding a definite answer, but unfortunately they haven't worked. For example, we can systematically compare and contrast the genes for the photosystems in bacteria, to build a gene tree that betrays the ancestry of the photosystems. Such trees are felled, though, by a fact of bacterial life – sex. Bacterial sex is not like our own, in which genes are inherited down the generations, giving rise to a nicely ordered family tree. Bacteria throw their genes around with a profligate disregard for the labours of geneticists. The result is more of a web than a tree, in which the genes of some bacteria end up in other, totally unrelated bacteria. And that means we have no real genetic evidence for how the photosystems came to be assembled together in the Z scheme.

But that doesn't mean we can't work out the answer. The great value of hypotheses in science is that, by making imaginative leaps into the unknown, they suggest new angles and experiments that can corroborate or refute the postulates. Here is one of the best – a beautiful idea from John Allen, professor of biochemistry at Queen Mary, University of London, and an inventive mind. Allen has the dubious distinction of being the one person I've written about in three consecutive books, with a different groundbreaking idea in each. Like the best ideas in science, this hypothesis has a simplicity that cuts straight through layers of complexity to the quick. It may not be right, for not all the great ideas in science are. But even if it's wrong, it shows how things *could* have come to be the way they are, and by suggesting experiments to test it, guides researchers in the right direction. It offers both insight and stimulus.

Many bacteria switch genes on and off in response to changes in their environment, says Allen; this in itself is common lore. One of the most important environmental switches is the presence or absence of raw materials. By and large, bacteria don't waste energy building new proteins to process raw materials if there aren't any around; they just close down the works until further notice. So Allen pictures a fluctuating environment – perhaps a stromatolite in shallow seas, in the vicinity of a hydrothermal vent that issues hydrogen sulphide into the world. The conditions would vary according to the tides, currents, time of year, hydrothermal activity, and so on. The critical factor is that Allen's hypothetical bacteria should possess both of the two photosystems, as cyanobacteria do today, but unlike cyanobacteria only ever use one of them at once. When hydrogen sulphide is present, the bacteria switch on Photosystem I and use it to produce organic matter from carbon dioxide. They can incorporate this new matter to grow, reproduce, and so on. But when conditions change, and the stromatolites are left without raw materials, these bacteria switch over to Photosystem II. Now they give up making new organic matter (they don't grow or reproduce any more), but they can maintain themselves by using sunlight to make ATP directly until better times. Each photosystem has its own benefit, and each evolved via a series of simple steps, as we've seen.

But what happens if a hydrothermal vent dies, or shifting currents lead to protracted changes in the environment? The bacteria must now rely on the

electron circuit of Photosystem II for most of the time. But here there is a potential problem: the circuit can bung up with electrons from the environment, even if it only happens slowly in the electron-poor surroundings. The electron circuit is a bit like a pass-the-parcel game. An electron carrier either has an electron or it doesn't, just as a child either has a parcel when the music stops, or does not. But now imagine a rogue supervisor with a pile of parcels; he keeps passing them into the children's circle one at a time. In the end all the children end up with a parcel each. Nobody can pass on their parcel; the game grinds to a halt amid confusion.

Much the same happens with Photosystem II. The problem is inherent in sunlight, especially in the days before an ozone layer, when more ultraviolet radiation penetrated down to sea level. Ultraviolet rays not only split water, but can also throw off electrons from metals and minerals dissolved in the oceans, first among them manganese and iron. And that would have presented exactly the kind of problem that stymied our pass-the-parcel game. A trickle of electrons enters the circuit.

Neither iron nor manganese is found in high concentration in seawater today, for the oceans are thoroughly oxidised; but in ancient times, both were plentiful. Manganese, for example, is found in massive amounts on the sea floor in the form of curious cone-shaped 'nodules', which accrete over millions of years around objects like sharks' teeth, one of the few bits of living things that can withstand the intense pressure at the bottom of the ocean. There are thought to be a trillion tonnes of manganese-rich nodules scattered across the sea floor, a huge but uneconomic reserve. Even the more economic reserves, like the massive Kalahari manganese fields in South Africa (another 13.5 billion tonnes of ore), were precipitated from the oceans, 2,400 million years ago. In short, the oceans were once full of manganese.

For bacteria, manganese is a valuable commodity: it works as an antioxidant that protects cells against the destructive power of ultraviolet radiation. When a manganese atom absorbs a photon of ultraviolet radiation, it throws off an electron, becoming photo-oxidised, and in the process 'neutralises' the ray. The manganese is 'sacrificed' instead of more important bits and pieces of the cell such as proteins and DNA, which otherwise would be shredded by rays; and so bacteria welcome manganese into their abode with open arms.

The trouble is that when these manganese atoms throw off an electron, it is always likely to be guzzled up by the 'grasping miser' form of chlorophyll in Photosystem II. And so the circuit gradually clogs up with electrons, just as our circle of children becomes swamped with parcels. Unless there is some way of bleeding off the excess electrons gumming up the circuit, Photosystem II becomes steadily less efficient.

How could bacteria bleed off electrons from Photosystem II? Here is the full genius of Allen's hypothesis. Photosystem II is clogged with electrons, while Photosystem I lies idle for lack of electrons. All that the bacteria need to do is disable the switch that prevents both photosystems from being active at once, either physiologically or by way of a single mutation. What happens then? Electrons enter Photosystem II from oxidised manganese atoms. They're then blasted up to a high energy level as the 'grasping miser' form of chlorophyll absorbs a ray of light. From here they're passed down the electron transport chain, using the release of energy to generate a little ATP. And then a diversion. Instead of returning to a gummed up Photosystem II, they are scavenged by the active Photosystem I, thirsty as it is for new electrons. Now the electrons are blasted up to a high energy level again, as the 'street hustler' form of chlorophyll absorbs a ray of light. And, of course, from here the electrons are finally passed to carbon dioxide, to generate new organic matter.

Does all this sound familiar? I have just described the Z scheme again. Just one single mutation connects the two photosystems in series, with electrons passing from manganese atoms, via the full Z scheme, to carbon dioxide to make sugars. What only now appeared to be an enormously convoluted and elaborate process is suddenly rendered virtually inevitable by a single mutation. The logic is flawless, the molecular parts are all in place and all serve a purpose as individual units. The environmental pressures are reasonable and predictable. Never did a single mutation make a bigger difference to the world!

It's worth a quick recap to appreciate the big picture in full relief. In the beginning there was a single photosystem, which probably used sunlight to extract electrons from hydrogen sulphide, and thrust them on to carbon dioxide to form sugars. At some point the gene became duplicated, perhaps in an ancestor of cyanobacteria. The two photosystems diverged under different usage.[9] Photosystem I carried on doing exactly what it had done before, while Photosystem II became specialised to generate ATP from sunlight by way of an electron circuit. The two photosystems were switched on and off according to the environment, but the pair were never switched on at the same time. Over time, however, Photosystem II has a problem, resulting from the properties of a circuit of electrons – any extra input of electrons from the environment jams up the circuit. It's likely there was a constant slow input of electrons from manganese atoms, used by bacteria to protect against ultraviolet radiation. One solution was to inactivate the switch, enabling both photosystems at once. Electrons would then flow from manganese, through both photosystems, to carbon dioxide, via a complex pathway that foreshadows the convoluted Z scheme in every eccentric detail.

We're now only one step short of full-blown oxygenic photosynthesis. We're drawing electrons from manganese, not from water. So how did the final shift occur? The surprising answer is that virtually nothing needed to change.

✦

The oxygen-evolving complex is the nutcracker that pinions water 'just so' to have its electrons cracked out one by one. When all the electrons are removed, the invaluable waste, oxygen, is flushed out into the world. The oxygen-evolving complex is really a component of Photosystem II, but sits at the very edge, facing the outside world, and gives a sense of being 'tacked on'. It's shockingly small. The complex is a cluster of four manganese atoms and a single calcium atom, all held together by a lattice of oxygen atoms. And that's that.

For some years, the irrepressible Mike Russell, whom we already met in Chapters 1 and 2, has argued that the structure of this complex is remarkably similar to some minerals cooked up in hydrothermal vents, such as hollandite or tunnel calcium manganite. But until 2006, we didn't know the structure of the manganese cluster in atomic resolution, and Russell's was a voice in the wilderness. But now we know. And although Russell was not quite right, his

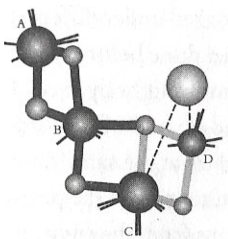

Figure 3.4 The ancient mineral structure of the oxygen-evolving complex – four manganese atoms (labelled A–D) linked by oxygen in a lattice, with a calcium atom nearby, as revealed by X-ray crystallography.

broad conception was absolutely correct. The structure, as revealed by a team headed by Vittal Yachandra at Berkeley, bears a striking resemblance to the mineral forms proposed by Russell (see Fig. 3.4).

Whether the original oxygen-evolving complex was simply a bit of mineral that got wedged in Photosystem II, we don't know. Perhaps manganese atoms became bound to oxygen in a lattice as they were oxidised by ultraviolet radiation, seeding the growth of a tiny crystal on site.[10] Perhaps the proximity of this cluster to chlorophyll, or to adjacent bits of protein, distorted it a little in some way, optimising its function. But whatever the origin of the cluster, there is a huge sense of the accidental about it. It is far too close to a mineral structure to be the product of biology. Like a few other metal clusters found at the heart of enzymes, it is almost certainly a throwback to the conditions found billions of years ago in a hydrothermal vent. Most precious of all jewels, the metal cluster was wrapped in a protein and held in trust for all eternity by the cyanobacteria.

However it formed, this little cluster of manganese atoms opened up a new world, not only for the bacteria that first trapped it, but for all life on our planet. Once it formed, this little cluster of atoms started to split water, the four oxidised manganese atoms combining their natural avidity to yank electrons from water, thereby releasing oxygen as waste. Stimulated by the steady oxidation of manganese by ultraviolet radiation, the splitting of water would have been slow at first. But as soon as the cluster became coupled to chlorophyll, electrons would have started to flow. Getting faster as chlorophyll became adapted to its task, water was sucked in, split open, its electrons drawn out, oxygen discarded. Once a trickle, ultimately a flood, this life-giving flow of electrons from water is behind all the exuberance of life on earth. We must thank it twice – once for being the ultimate source of all our food, and then

again for all the oxygen we need to burn up that food to stay alive.

It's also the key to the world's energy crisis. We have no need for two photosystems, for we're not interested in making organic matter. We only need the two products released from water: oxygen and hydrogen. Reacting them together again generates all the energy we'll ever need, and the only waste is water. In other words, with this little manganese cluster, we can use the sun's energy to split water, and then react the products back together again to regenerate water – the hydrogen economy. No more pollution, no more fossil fuels, no more carbon footprints, no more anthropogenic global warming, albeit still some danger of explosions. If this little cluster of atoms changed the makeup of the world long ago, knowing its structure should be the first step to changing our own world today. As I write, chemists around the world are racing to synthesise this tiny manganese cluster in the lab, or something similar that works as well. Soon, surely, they will succeed. And then it can't be long before we learn to live on water and a splash of sunshine.