LIAM GRAHAM

MOLECULAR STORMS

THE PHYSICS OF STARS, CELLS AND THE ORIGIN OF LIFE

Springer

Molecular Storms

Liam Graham

Molecular Storms

The Physics of Stars, Cells and the Origin of Life

Liam Graham ⓘD
London, UK

Cover image: Yukhym Turkin

Advance Praise for This Book

"Only a few writers have managed to turn the highly technical jargon of science into language accessible for interested lay readers. Isaac Asimov showed us how it could be done, and Carl Zimmer and Brian Greene are continuing today. In Molecular Storms, his first book, Liam Graham has shown that he has the essential quality required to join this group, a love of first learning then explaining how the universe works."

David Deamer, Professor of Biomolecular Engineering, University of California, Santa Cruz, author of *Assembling Life*

"Following in the footsteps of Stephen Hawking's 'A brief history of time' and Simon Singh's 'Fermat's Last Theorem' this exceptionally accessible book will leave you marvelling at the wonders of the world and, if you didn't listen to your science teachers, wishing you had. Graham writes with the mind of a physicist and the soul of a poet."

Nicki Hayes, CCO, The Communication Practice, author of *First Aid for Feelings*

"A great place to start if you are interested in the origin of life! Graham skillfully presents an up-to-date account of some of the deepest problems in science: how does living matter work, and how might it have come to be? With clear explanations, simply presented and intuitively appealing, the level is accessible to a curious non-expert yet still a profitable introductory read for an undergraduate or interdisciplinary researcher in life or physical sciences."

Nigel Goldenfeld, Professor of Physics, University of California, San Diego

"This book should be called 'Everything You Always Wanted To Know About Thermodynamics (But Were Afraid To Ask)'. Through humour and powerful images, it took me on an enlightening and enjoyable guided tour through a hitherto obscure corner of physics. Reading it changed the way I think about the world and my reasons for being in it."

Valentine Allen, actor, director, author of *Rondeurs*

"I'm delighted to see thermodynamics portrayed as a source of 'joy' and I'd have appreciated finding this book as a senior in college."

Nicole Yunger Halpern, Professor of Physics, University of Maryland, author of *Quantum Steampunk*

"Liam Graham has that rare ability to take a seemingly staggeringly complex subject—life itself—and strip it down into a set of smaller problems. He then shows you, by argument, analogy, vivid examples and good-humoured persuasion, that each of these smaller problems is much more manageable, and can be related back to a few key concepts. Taken one at a time, each then becomes a staging point in your journey towards understanding the big picture. Like all the best guides, his enthusiasm and sense of wonder are infectious."

Stephen Wright, Professor of Economics, Birkbeck College, University of London

…if we were to name the most powerful assumption of all, which leads one on and on in an attempt to understand life, it is that all things are made of atoms and that everything that living things do can be understood in terms of the jigglings and wigglings of atoms.

—Richard Feynman

In memory of Gillian Rose (1947–1995)

Preface

Can physics explain life? Living cells, like everything else, are subject to the laws of physics. Yet take a look inside a cell and you are met with staggering complexity. A network of interlinked structures, many of them molecular machines with moving parts, perform a complex dance of minutely choreographed activity. This is unlike anything in our experience. The culmination of this activity is something equally unique: the cell produces a copy of itself.

The aim of this book is to uncover the physics behind this complexity. There are two intertwined stories. One is about how cells work. This is a story across scales from atoms to molecules to the systems of complex molecules we call living cells. The other is about the conditions necessary for life to exist, a story across time, from the start of the universe to stars and planets to the molecular playgrounds where life started.

Part I of the book presents the scientific concepts that underlie these two stories. Part II uses these concepts to explore the prerequisites for life. Applied to the universe as a whole, we'll see how physics explains the existence of planets capable of supporting life. Applied to chemistry, we'll discover systems which hint at the first steps from non-living to living matter. Part III uses all of this to unravel the workings of a living cell. It starts with the strange nature of the environment inside cells, moves on through cellular processes and the machines which control them, then turns to cell division and reproduction. Finally, we'll see how physics can clarify the biggest question of all: how life began.

Two notes of humility are in order. First, our understanding of living systems is far from complete and fascinating open questions lurk everywhere. Second, none of the research I discuss is my own. My contribution is to weave together material from across disciplines, choosing work that fuels my decades-long fascination with these matters.

This book is intended for all those who are curious about how the world works. I hope that scholars, students, and general readers will appreciate a non-technical overview of some of the most exciting ideas in science. Specialists in any of the fields I cover may find my treatment superficial. If this is the case, I hope that my book still conveys the most remarkable thing of all: that science has progressed to a point where this exercise is possible.

London, UK   Liam Graham

Acknowledgments

Contents

About the Author

Liam Graham "Do we need more than physics to understand the world?" Liam first asked himself this question as a teenager and it has been the driving force behind his career ever since. After a degree in Theoretical Physics at Cambridge and a master's in Philosophy at Warwick, he eventually found economics to be an appealing middle ground and completed a Ph.D. at Birkbeck College, London. To pay the rent, he taught English, developed and sold trading software and was the numbers' guru for a boutique finance house.

Liam's 15 year academic career was mostly spent as an Associate Professor at University College London, working in one of Europe's top Economics departments. His research involved building mathematical models of an extremely complex system, the macroeconomy, and his work was published in all the top macroeconomics journals. Whether working on philosophy or economics, he never stopped reading science and exchanging with scientists. In 2018 he left UCL to concentrate on his original question and the wide-ranging, multidisciplinary and endlessly fascinating project it has become.

1

Introduction

If you want to understand the world atom by atom, you need quantum physics. If you happen to be near a black hole, you'll need general relativity as well. These two pillars of twentieth century physics describe the world with exquisite accuracy. But neither are much use for many things that may interest non-physicists: large molecules like DNA; the assemblies of such molecules we call living cells; planetary oceans and atmospheres; human technology; or even a glass smashing on the floor. To understand these, you need the third pillar of modern physics, the poor cousin, thermodynamics.

Thermo-dynamics. The study of heat in motion. The name has stuck even though it conjures up images of billowing steam and bearded engineers. Thermodynamics was indeed a child of the industrial revolution. Then, at the end of the nineteenth century, it was given statistical foundations. This turned it into something much more powerful: a way of describing things that are made from too many parts to keep track of. A way of describing almost everything.

Simple systems can be modelled directly. The computing power of your smartphone and some high school physics are enough to calculate the trajectory of a probe as it moves through the solar system. Broadly the same physics applies to the motions of molecules. But in a bacterium there are 10 billion molecules. In a litre of air there are 10 billion trillion. However big a computer you might have, you'll never be able to keep track of them let alone calculate their interactions. This is where thermodynamics comes in.

One of the peculiar things about thermodynamics is that no physics is required to understand some of its central concepts: exploration, equilibrium, entropy and the second law. They are statistical properties of systems made up of parts whose configuration changes randomly. I'll introduce them by asking you to picture a grid of flashing lights. It's hard to overstate how odd this is: we can get to a fundamental law of nature by studying random patterns.

Randomness is right at the heart of thermodynamics and underlies the second law itself. In a physical system, the randomness arises from the incessant motion of molecules. Imagine pollen grains suspended in water. If you observe them closely, you'll see that they wiggle around constantly and unpredictably. This wiggling can be seen under a low-power microscope and was first recorded in the nineteenth century. A few decades later, one of the papers that made 1905 Einstein's annus mirabilis explained the motion as resulting from the constant bombardment of the pollen grains by water molecules. I'm going to call this bombardment the molecular storm. It is central to everything that follows.

The microscopic details of systems constantly fluctuate as the storm drives them to explore patterns. An equilibrium is when this randomness washes out and, viewed as a whole, the system is unchanging. Think of the air around you. While its molecules are in constant movement, its pressure and temperature are constant. If a system is not already in equilibrium, the storm will drive it there. Reaching equilibrium may take milliseconds or almost forever.

Equilibrium itself is (mostly) dull. Along the path to equilibrium, interesting things can happen. Engines are devices which exploit systems on their way to equilibrium. Examples are the plughole vortex in your bathtub; the Great Red Spot of Jupiter; a waterwheel; the motor of your car; or the nanoscale engines that propel a bacterium. Some engines self-organise. Some are built by other processes. Some involve a demon, either metaphysical, mechanical or chemical.

Thermodynamics may be the poor cousin of modern physics, but it is also surprisingly versatile. With a few concepts and a few laws, we can gain insights into more philosophical questions: the nature of time, the initial and final states of the cosmos and the possibility that you, gentle reader, may be a random fluctuation in a long-dead universe.

That is the first part of the book, introducing the basic concepts that will structure the remainder. The second part applies them to explore the prerequisites for living systems. While thermodynamics can tell us some general things about the start of the universe, cosmology supplies the details. Soon after the big bang the universe was a very hot cloud of particles and radiation. Thermodynamics then explains how this apparently unpromising initial

state leads to the universe we live in, teeming with rich structure at every scale from galaxies to stars to planets. And how the surfaces of these planets become molecular playgrounds, full of possibilities for interesting things to happen.

To understand these molecular playgrounds we need chemistry. Chemical terminology—reactions, chemical equilibrium, catalysts—conveniently bundles up quantum physics with more thermodynamics. The language of chemistry is essential to describe living systems. Also, we will see that surprisingly simple chemical systems can show analogues of properties we usually associate with life: growth, reproduction and regulation. The molecular storm drives all of this.

The third part of the book turns to living cells. Once again, the molecular storm is behind everything. It drives proteins to fold and structures to self-assemble. It drives food molecules to where they can be used and distributes their products around the cell. This happens in an environment far from anything in our experience. Unlike laboratory test tubes, living cells are jam-packed with large molecules. It is in this crowded environment that the molecular storm rages. Scaled up to human dimensions, it would be a wind of tens of thousands of kilometres per hour, blowing from constantly changing directions, in air thicker than treacle.

Another difference from test-tube chemistry is that many of the large molecules in cells are nanomachines, molecular devices that work by mechanical movements. There are gates which close to trap some molecules and expel others. There are rotary motors driven by protons. Proton pumps driven by electrons. Nanoengines which build proteins, amino acid by amino acid, in precisely regulated sequences. Nanoengines that replicate DNA with astonishing accuracy and create the structures necessary for cells to divide. All of these are driven by the molecular storm and thermodynamics is essential to understand how they work. With this understanding, we can then investigate the origin of life and the possibilities for life elsewhere in the universe.

These are the broad outlines of the story I'm going to tell. A few minutes after the big bang, the temperature of the universe was around 10 billion °C. In the unthinkably distant future, when the last black hole evaporates, the temperature of the universe will be an infinitesimal fraction above absolute zero. Everything around us, the wild complexity of the world, life on earth and we ourselves are a fleeting stage on this long path to equilibrium. To make sense of all this, you need thermodynamics.

Part I

Molecules in Motion

2

Random Patterns

Summary Some of the central concepts of thermodynamics can be introduced without using any physics. I'm going to ask you to imagine a device consisting of lights changing colour randomly. This will let us examine the ideas that structure the rest of the book: microstates and macrostates, fluctuations, exploration and equilibrium. Then comes a definition of entropy, a statement of the second law and a discussion of the arrow of time. The chapter ends by pointing out the limitations of our intuitions about order and disorder.

Thermodynamics is a contender to be the worst named part of physics. Entropy is surely one of its least understood concepts. Don't take my word for this. John von Neumann was a colossus of twentieth century science, making fundamental contributions to fields including maths, quantum physics, fluid mechanics, game theory, the theory of computation and the cold-war policy of mutually assured destruction. In the late 1940s, Claude Shannon came up with a new way of describing information. He asked von Neumann what he should call it. Von Neumann replied:

> You should call it entropy, for two reasons. In the first place your uncertainty function has been used in statistical mechanics under that name, so it already

has a name. In the second place, and more important, no one knows what entropy really is, so in a debate you will always have the advantage.[1]

Although we may not know what entropy really is, we have an intuitive concept that in many situations is a good-enough approximation. A glass on a table is low entropy. A smashed glass on the floor is high entropy. A tidy room is low entropy. A messy room is high entropy. To get from a state of high entropy to a state of low entropy requires effort: tidying the room, putting the fragments together again.

This chapter will introduce the basic concepts of thermodynamics, including entropy and the second law, concepts which are central to the book's argument. At the end of the chapter I'll return to the intuitive concept of entropy and discuss its limitations. I'll do all this using a simple model of flashing lights. If you prefer thinking in terms of a real physical system, I suggest you jump to the next chapter and refer back to this one if you need to.

2.1   A Beginner's Guide

Science is about taking measurements of physical systems, the temperature of a liquid, the pressure of a gas, the colour of a surface, how well ordered a bookcase is, how untidy a room is. These are measurements of the overall state of the system. All these systems are made up of components: atoms, books or whatever is in the room. What do aggregate measurements tell us about the state of the components? This is the question that entropy answers.

As von Neumann's comment suggests, entropy is a concept that makes people nervous. To try to ease such apprehensions, here's a very brief example. Take two dice. Ask a friend to throw them and tell you only their total. What does this total tell you about the values of the individual dice?

Let's call the numbers on the two dice their state. Entropy is a count of the number of states corresponding to each total. If the total is 2, then only a single state (1,1) is possible, so entropy is 1. If the total is 4, three states (1,3; 3,1; 2,2) are possible, so entropy is 3. If the total is 7, there are 6 possible states (1,6; 6,1; 2,5; 5,2; 3,4; 4,3), so entropy is 6. Entropy is a way of counting.

The higher is entropy, the greater is the uncertainty about the state of the individual dice. If entropy is 1, the total tells you the state of the dice and there is no uncertainty. If entropy is 6, the total tells you that the dice could

---

[1] Tribus and McIrvine (1971).

be in one of 6 possible states. To look at it the other way round, the entropy of a total is also a measure of its probability. Shake the dice, and you are most likely to obtain a total corresponding to maximum entropy.

This is the essence of thermodynamics. Take a system made up of components that are randomly rearranged. Take a measurement of the system as a whole. Randomise the system, and it will tend to end up at the measurement corresponding to maximum entropy. But there are not many dice in nature, so the rest of the chapter will use a different model.

2.2   Red and Blue Lights

Let's invent a device. It consists of a rectangular grid of lights, which can be either red or blue. Each light is attached to a circuit which, when activated, randomises the colour. When the circuit is activated the light has an equal probability of being red or blue, just like flipping a coin. The grid is attached to a clock and, at each tick, a fixed number of the lights, chosen randomly, are flipped. There is a readout which measures the proportion of lights that are red. Once again, we have a system made up of components and we take measurements of the system as a whole.

Picture 100 lights, arranged in a 10×10 grid. Let's say that every second 10 of these lights are flipped. At the start, all of the lights are red. Then switch on the device. Immediately, patches of blue start appearing in the sea of red. As time passes, some of these blue patches turn back to red and some areas of red turn blue. After a while, around half the lights are red and around half blue but with a pattern that changes with every tick of the clock. The readout starts at 100% (all lights red) then gradually falls to around 50% (half of the lights red) and then fluctuates around this value, sometimes higher, sometimes lower.[2]

Now let's look at things in a bit more detail. Start with just three lights, all red giving a measurement of 100%. Then leave the device running for a while. If we then take another measurement, it will be one of four values: 0, 33.3, 66.7 or 100%. This is shown in Table 2.1. There is a 1/8 chance that the system will be in the same state (all red) as it started. There is a 1/8 chance the lights will all be blue. If we leave the machine running and take regular measurements their average will be 50%.[3]

---

[2] For a video, see www.TheMaterialWorld.net.
[3] The average is $(1 \times 0 + 3 \times 33.3\% + 3 \times 66.7\% + 1 \times 100\%)/8 = 50\%$.

Table 2.1 Patterns of 3 lights

States	Number of states	Proportion of red (%)
BBB	1	0
BBR, BRB, RBB	3	33.3
RRB, RBR, BRR	3	66.7
RRR	1	100

Table 2.2 Patterns of 5 lights

States	Number of states	Proportion of red (%)
BBBBB	1	0
BBBBR, BBBRB, BBRBB, BRBBB, RBBBB	5	20
BBBRR etc.	10	40
BBRRR etc.	10	60
RRRRB, RRRBR, RRBRR, RBRRR, BRRRR	5	80
RRRRR	1	100

How does this behaviour change with the number of lights? For 5 lights, there are 6 possible measurements and 32 states as shown in Table 2.2. After the device has been running a while, the probability that all red will be observed is now 1/32. If we keep taking measurements they will be between 40 and 60% around two-thirds of the time and their average will again be 50%.[4]

Box 2.1 Numbers large and small

I'm going to need to refer to both very small and very large numbers. Standard notation goes in powers of 10: a hundred is 10^2, a thousand 10^3, a million 10^6 and a billion 10^9. Put a minus sign in front of the power and add "th" to the word; 10^{-2} is a hundredth.

It's useful to have some points of comparison.[5] We could use the Planck time, 10^{-44}s, the size of an atom, 10^{-10}m, the number of seconds since the big bang, 10^{17}, or the number of atoms in the universe, 10^{79}.

Thermodynamics involves thinking about systems made up of components. As the number of components increases, the number of arrangements rapidly explodes. We don't need to go as far as simians and Hamlet. This box contains around 1,000 characters, just one among 26^{1000}, or 10^{1400}, combinations. Even

---

[4] The average is $(1 \times 0 + 5 \times 20\% + 10 \times 40\% + 10 \times 60\% + 5 \times 80\% + 1 \times 100\%)/32 = 50\%$.

[5] The calculations are in Appendix A.1.

if we think of one for every atom in the universe and every tick of the Planck clock since the big bang, we don't get close. And this is in a system made up of just 1,000 parts. Such numbers defy comparisons. To refer to them, I'm going to adopt the philosopher Dan Dennett's term, VAST, for "Very much more than ASTronomical".[6]

If you want to make even these numbers shrink into insignificance, you need something really large, such as Graham's (no relative) number[7] or infinity, more of which in Chap. 4.

As we increase the number of lights, three things happen

- The number of patterns explodes.[8] For 3 lights, there are 8 patterns, for 5 there are 32, for 15 there are around 33,000, for 30 there are over a billion and for 100 there are around 10^{30}.
- An increasing proportion of these patterns corresponds to a measurement close to the average of 50%.
- The probability of returning to all red gets smaller.

For more lights the magnitudes get out of control. For 260 lights there are 10^{79} combinations, one for each atom in the universe. We're approaching the realm of VAST numbers, see Box 2.1.

To avoid having to deal with such large numbers, let's instead concentrate on the aggregate measurement. Table 2.3 shows the percentage of patterns with measurements within a certain distance of the average. To pick out a couple of examples, with 100 lights, 68% of measurements will be between 40 and 60%. With 100,000 (10^5) lights 99.8% of measurements will be between 49 and 51%. With a billion (10^9) lights, 99.8% will be between 49.99 and 50.01%. The final entry in the table tells us that, with 10 billion lights (10^{10}), every time a measurement is taken it will give the result of 50% to an accuracy of 0.01%. 10^{10} is a lot of lights, but it is tiny compared to the 10^{22} molecules in a litre of air.

Let's now use our device to introduce some terms and concepts which will be used throughout the rest of this book. The next chapter returns to them in the context of a real physical system.

---

[6] Dennett (2013).

[7] Gardner (2001), Chap. 33.

[8] The number of combinations for n lights is 2^n. A rule of thumb is that $2^n \approx 10^{n/3.3}$ so $2^{100} \approx 10^{30}$.

Table 2.3 Measurement

Number of lights	Probability of measurement within 10% of average (%)	Probability of measurement within 1% of average (%)	Probability of measurement within 0.01% of average (%)
100	68	0.1	0
1,000	99.8	24	0
10,000	100	68	0
100,000	100	99.8	0
10^6	100	100	0.1
10^7	100	100	24
10^8	100	100	68
10^9	100	100	99.8
10^{10}	100	100	100

2.3   Microstates and Macrostates

The *microstate* of a system is the state of its components. For our device, this means the pattern of lights. The first columns of Tables 2.1 and 2.2 show the microstates.

The *macrostate* of a system is some measurement of the system as a whole, an aggregate measurement, an observable. For our device, this is the proportion of lights that are red, the third column of Tables 2.1 and 2.2.

Each microstate corresponds to a single macrostate. But one macrostate may correspond to many microstates. Table 2.2 shows that with 5 lights, a macrostate of 20% could arise from 5 different microstates or a macrostate of 40% from 10 different microstates.

These concepts can be applied to any physical system made up of smaller components. For a box full of air, the macrostate is what we choose to measure—temperature or pressure, for example—and the microstate relates to the details of its component molecules. For a room full of objects, the macrostate is some measure of tidiness, the microstate the position of the individual objects.

2.4   Exploration

The mechanism of the device randomises the colour of its lights. With each tick of the clock, the pattern or microstate of the system changes. Or in other words, the device *explores* different microstates.

Will it explore all microstates? This will depend on how many lights there are, how many lights are flipped with each tick of the clock and how long the