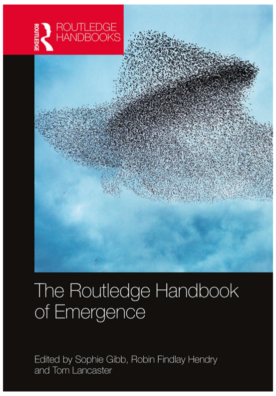

The Routledge Handbook of Emergence

Sophie Gibb, Robin Findlay Hendry, Tom Lancaster

Phase Transitions, Broken Symmetry and the Renormalization Group

Publication details
https://www.routledgehandbooks.com/doi/10.4324/9781315675213-20
Stephen J. Blundell
Published online on: 22 Mar 2019

19

PHASE TRANSITIONS, BROKEN SYMMETRY AND THE RENORMALIZATION GROUP

Stephen J. Blundell

Introduction

A piston is pushed into a cylinder and compresses a gas. The pressure goes up as the volume goes down, in accordance with Boyle's law. In such a simple thermodynamic description, it is not necessary to worry about what is going on at the atomic scale. Nor do we need to concern ourselves with the fact that this simple lab experiment is being carried out on the surface of planet Earth, a 12,800-km-diameter rotating sphere filled with iron and coated with a silicate crust. Nor is our Boyle's law analysis using the fact that such a rotating sphere glides along its orbit through largely empty space. Nor indeed do we need to keep in mind that the galaxy in which the solar system resides is an unusually high-density region of the universe. Just as our thermodynamics experiment can be performed in a lab without fussing over such cosmological niceties, a cosmologist in the office next door who is studying the Friedmann model of the expansion of the universe will take the average mass density of the universe (equivalent to around a few protons per cubic metre) as a simple parameter. She will blissfully ignore the "clumpiness" of matter at smaller scales, and of course will neglect the presence of the experiment with the piston and the cylinder of gas sitting on planet Earth.

Thus we take it as a given that when a portion of the universe is selected for study, be it a gas or a galaxy, we are allowed to blissfully ignore what is going on at scales that are much larger, or indeed much smaller, than the one we are considering. We might perhaps use those different scales to provide appropriate matching conditions. For example, the fact that the experiment with the piston is located on the surface of planet Earth would inevitably introduce a gravitational field into our experiment, which might perhaps be relevant. By considering the atomic-level motion of the gases we would be able to match a kinetic-theory description containing atomic masses and velocities on to our thermodynamic macroscopic description, but we could also ignore the existence of atoms entirely (as most scientists did in the nineteenth century) and still perform perfectly acceptable thermodynamics. These other scales, the atomic or the planetary, only provide at most the parameters that determine the boundary conditions, and the details of physics arising at other scales are neglected.

However, it turns out that there is a class of very familiar problems where this simple neglect of what is going on at other length scales fails. In these problems it is not only important, but also unavoidable, to consider the physics on multiple length scales that can vary over many orders of magnitude. A very good example of such a problem is a phase transition. In fact, you may become aware of this issue whenever you boil the water in a kettle to make yourself a cup of tea, if you pay attention and listen carefully. As the water heats up, its temperature increases. There is a simple relation between the amount of energy you put into the water per second (the power of the kettle, probably written on a manufacturer's label attached to the base) and the temperature rise per second. Initially this process of warming proceeds steadily and quietly (save a brief hissing noise you hear shortly after switching the kettle on, as dissolved gases are driven out of the water). However, as the heating process nears completion, the kettle starts to become rather noisy due to the hot water bubbling and frothing, and the sound gets progressively louder. When the temperature hits 100 degrees Celsius, the liquid water does something remarkable. It begins to change into a gas and steam starts to rise upwards from the spout. This is an example of a phase transition, a change of state of water. It is utterly discontinuous and strongly contrasts with the smooth continuous changes that have occurred previously (42-degree water is similar, but just a bit hotter, than 41-degree water, which is similar, but a bit hotter still, than 40-degree water, etc.; conversely, steam at 100 degrees Celsius is radically different from water at 100 degrees Celsius). It also involves a finite slug of energy supplied from the kettle to make it happen; this is known as latent heat and indicates that the two phases have different entropy.

But not only is the transition remarkable in its discontinuity, the region close to the transition is remarkable too. Fluctuations occur near a phase transition and are produced on all length scales. Thus, not only does water close to the boil contain lots of little bubbles of steam, but there are big bubbles, and occasionally very big bubbles. This gives rise to the noisy, boisterous behaviour of hot water close to the boil. Bringing water to the boil in a pan on a stove demonstrates this even more vividly, but the same effect is apparent already in the kettle.

We can obtain a visual representation of this behaviour near a phase transition by examining snapshots of configurations of a magnetic system (a two-dimensional Ising model, which I will explain later) that also exhibits a change of state at a critical temperature T_c at which it changes from a ferromagnetic state ($T < T_c$) to a paramagnetic state ($T > T_c$). The snapshots are shown in Figure 19.1 and illustrate the magnetic configurations after the system has had time to settle down at different temperatures. Each picture has regions that are either black or white, corresponding to atomic magnetic moments (spins) that are only allowed to be up or down. At low temperature all the spins are aligned and the region is uniformly coloured black, representing identical spin states pointing up. As the temperature increases, a few white blobs begin to appear as small fluctuations start to occur more readily. Close to the transition, the fluctuating regions are very large and the system starts to lose its net magnetic moment. Above the transition, fluctuations occur more readily, but the size of correlated regions of spins (known as the correlation length) begins to decrease and, at the highest temperature shown, the state of each individual spin is only weakly correlated with its neighbour. At temperatures close to T_c, the fluctuations become very slow, due to the correlation length of fluctuating spins becoming very large. This is because to cause a fluctuation, it is necessary to flip a large, correlated block of spins, which is a slow process. Both the correlation length and the fluctuation time diverge at T_c (the latter phenomenon is known as "critical slowing down"). This slowing down affects the Monte Carlo simulations that were used to generate Figure 19.1 just as much as in the real system (in both cases, interactions between more distant spins are becoming important; both the real system and the simulated system become more correlated and complicated near T_c), and this increases the convergence time of calculations near T_c. In general, what we call critical behaviour (a class of phenomena observed at temperatures close to phase transitions) is characterized by the onset of long-range correlations. A diverging correlation length leads to singularities in response functions, and at T_c the system becomes scale-invariant. Techniques such as the renormalization group, which I will describe

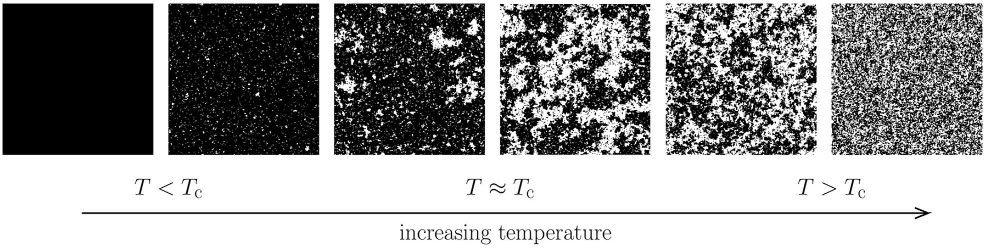

Figure 19.1 A Monte Carlo simulation of spins in the two-dimensional Ising model at various temperatures. Spins are coloured black or white according to whether they are in the up or down state, respectively.

[From Blundell and Blundell (2010)]

later, exploit the scale invariance. It is found that length scaling (Wilson and Fisher 1972; Wilson 1972) accounts for the universality of behaviour found in systems with very different underlying physics because many of the system parameters turn out not to be important at the transition (in the language of the renormalization group, they are said to be "scaled away") and only a very few relevant parameters (such as the system and order parameter dimensionality) remain.

Phase transitions are thus remarkable phenomena in which it is not possible to work at one scale and ignore all the others. You have to think about all scales of length. Counterintuitive behaviour can be expected in such scale-invariant systems, and there are important analogies with self-similar geometric structures, such as fractals (Stinchcombe 1989), that also possess scale invariance.

It is worth recalling that the "usual" scaling behaviour is at the heart of much familiar physics. This use of calculus in physics, due to Newton and Leibniz, relies on the notion that scaling is simple and limits behave sensibly. Thus, quantities such as velocity v can be evaluated by considering a small displacement Δx in a small time interval Δt and writing $v = \Delta x/\Delta t$. This formula is taken in the limit when Δt tends to zero and allows differential equations to be constructed. For the usual $x(t)$ that describes planetary orbits and the motion of cannonballs, "taking the limit" works because, on a sufficiently short time scale, $x(t)$ becomes linear. Thus, halving Δt means halving Δx. The trouble occurs in systems for which these limits do not behave so straightforwardly in the limit of the very smallest distance and timescales. Consider an electron with charge Q. If you model it as a little ball of radius R, with the charge Q smeared uniformly over its surface, then the electric field energy stored in the charge is given by $Q^2/(8\pi\varepsilon_0 R)$. This stored energy is equivalent to a mass, so the particle mass m is then equal to $m = m_{\text{bare}} + Q^2/(8\pi\varepsilon_0 Rc^2)$, where c is the speed of light and m_{bare} is the "bare mass", the mass the particle would have without the effect of the electric field energy. This means that the mass m that we measure in the experiment is not the intrinsic mass of the particle but contains this additional contribution. But worse, if $R = 0$ (and we think the electron "is" a point-like particle), then this electric field energy contribution is infinite! A similar effect complicates our understanding of the electric charge of the electron, since during very short time intervals vacuum fluctuations result in the creation (and very quickly afterwards, the annihilation) of particle–antiparticle pairs and these pairs will screen the electric charge. This means that the "bare charge" of the particle will appear to be reduced when viewed from a distance R, and in fact the apparent charge Q will diverge as R is sent to zero. This worrying state of affairs has been addressed in quantum field theory by renormalization techniques (Delamotte 2004; Weinberg 1995). I will not review these ideas here insofar as they have been applied to the problem of the mass and charge of elementary particles, but rather concentrate on how they can illuminate the problem of phase transitions. First, it is necessary to consider the effect of symmetry breaking.

Symmetry breaking

As a liquid cools there can be a very slight lowering of the volume, but it retains a very high degree of symmetry. However, below the melting temperature, the liquid becomes a solid and that symmetry is broken. This may at first sight seem surprising because a picture of a solid with all the atoms lined up in a regular lattice "looks" more symmetrical than that of the liquid in which the atoms are floating around all over the place. The crucial observation is that any point in a liquid is, on average, exactly the same as any other. If you average the system over time, atoms visit each position in space with uniform probability. There are no unique directions or axes along which atoms line up. In short, the system possesses complete translational and rotational symmetry. In the solid, however, this high degree of symmetry is nearly entirely lost. Rather than being invariant under arbitrary rotations, the lattice of atoms might be invariant under, say, four-fold rotations; rather than being invariant under arbitrary translations, it is now invariant under a translation of an integer combination of lattice basis vectors. Therefore, not all symmetry has been lost but the high symmetry of the liquid state has been, to use the technical term, "broken". It is impossible to change symmetry gradually. Either a particular symmetry is present or it is not. Hence, phase transitions are sharp and there is a clear delineation between the ordered and disordered states.

Not all phase transitions involve a change of symmetry, and one such example is the liquid–gas transition, illustrated in the phase diagram in Figure 19.2 that shows the three phases: solid,

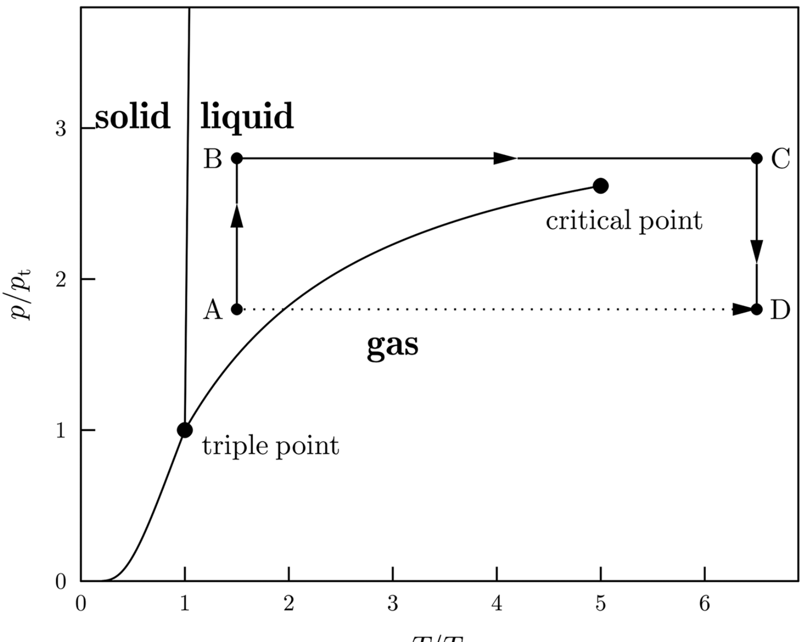

Figure 19.2 The phase diagram of a substance, showing how one can avoid the phase transition by going around the critical endpoint (route ABCD) rather than straight across the phase boundary (route AD)

liquid and gas. There really is no fundamental difference between a very dense gas and a liquid. Thus, the boundary line between the liquid and gas regions is terminated by a critical point. This means that it is possible to "cheat" the sharp phase transition by taking a circuitous path (ABCD) through the phase diagram that avoids a discontinuous change because you never cross the phase boundary; you go around it. For example, you can start with a liquid at low pressure and temperature (configuration A), pressurize it (B), then warm (C) and finally release the pressure (D). Configuration A is firmly in the liquid region, while configuration D is firmly in the gaseous region, but the route from A to D has avoided the phase transition. Of course, if you take the system straight from A to D by warming at constant pressure (exactly as you do when you boil the water in a kettle), then you observe a sharp discontinuity in properties as you cross the phase boundary. The fact that you can, if you want, *avoid* the transition by a judicious choice of route through the phase diagram is a consequence of the two phases having the same symmetry.

The same trick cannot be played when you transform a solid into a liquid because the two phases have different symmetries, and symmetry is not a property that you can have in anything other than full measure. To restate, either the system possesses a particular symmetry or it does not, and so symmetry-breaking transitions cannot be smooth crossovers. Consequently, there is no critical point for the melting curve and no way to cheat. The set of symmetry-breaking phase transitions includes as members those between the ferromagnetic and paramagnetic states (in which the low-temperature state does not possess the rotational symmetries of the high-temperature state) and those between the superconducting and normal metal states of certain materials (in which the low-temperature state does not possess the same symmetry in the phase of the wavefunction as the high-temperature state).

The role of system size

To understand some of the conceptual issues underlying symmetry breaking, it is helpful to consider a very simple model. Take N^2 spins (magnetic moments) arranged on a lattice of $N \times N$ sites. We will start with four spins on a 2×2 lattice. Now allow each spin to take one of only two possible configurations: up and down. Now add an interaction between the spins, operating only between nearest neighbours, and which saves energy J if the spins are aligned and costs energy J if the spins are antialigned. The 16 possible configurations are shown in Figure 19.3. This is known as the Ising model (which is usually studied on a much bigger lattice). If we take $J > 0$, then the configurations with the aligned spins (the two ones on the left) are the lowest in energy.

We might expect to find magnetic order in our simple model, but as we shall see, it does not emerge as simply as one might at first think. In statistical mechanics, we would usually

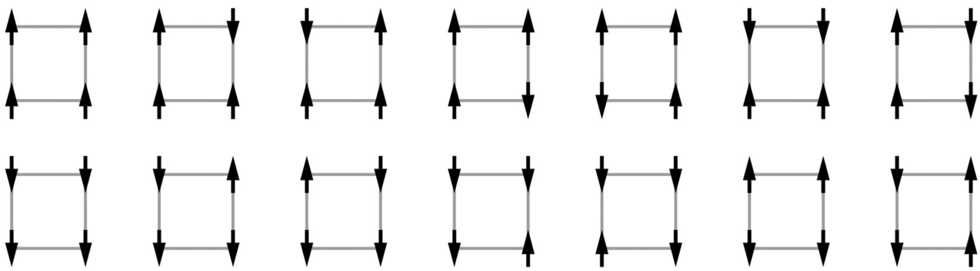

Figure 19.3 The 16 configurations of the Ising model on a 2×2 grid. Each configuration is shown with their symmetry-reversed partner above or below.

assume that if our system were connected to a thermal reservoir (a heat bath) at temperature T, then on average it would adopt a configuration given by a thermal average. The two easy limiting cases to consider are $T = 0$ (when there is no energy to be borrowed from the thermal reservoir and only the lowest energy configuration can be occupied) and $T = \infty$ (when energy can easily be exchanged with the thermal reservoir and every configuration is occupied with equal likelihood). Note that the problem is insensitive to which direction we call up. Thus, at $T = 0$, for example, the two configurations with (i) all spins up and (ii) all spins down (the two ones on the left of Figure 19.3) both minimize the energy and are therefore occupied with probability ½. The expected total number of up-spins minus the down-spins (the net magnetization) is therefore zero (because we weight the two configurations the same and they cancel out). The magnetization is also zero at $T = \infty$ because all configurations are now occupied and there is no net preference for up over down.

In fact, this is true at all temperatures as we can easily show. At a general temperature, the probabilities will be proportional to the Boltzmann factors $\exp(-E/k_BT)$, where E is the energy of a particular configuration. The expected magnetization can be written as $\langle M\rangle = \sum_i M_i \exp(-E_i/k_BT)/\sum_i \exp(-E_i/k_BT)$, where E_i and M_i are the energies and moments in the ith configuration. $\langle M\rangle$ will still be zero at any temperature because each term in the sum, which corresponds to a particular configuration, has a corresponding term with opposite magnetization, corresponding to a configuration in which all the spins are reversed (the configurations in Figure 19.3 have been arranged so that the symmetry-reversed configurations are grouped together, one drawn above the other). Thus, statistical mechanics, on its own, does not give us the non-zero order that we are looking for. This should not be surprising because the underlying model (in this case our Ising model) is entirely symmetrical between the up-spins and the down-spins; therefore, the solutions to this model possess this symmetrical property and there is nothing to select the up-spins over the down-spins. What is needed is some kind of symmetry *breaking*, where the system chooses one particular configuration and ignores its symmetrically related partner.

Experiment shows magnetic order does occur, so where does it come from? We can try studying the Ising model on a computer to learn more about it. A magnetic system modelled in a Monte Carlo simulation adopts a particular configuration and, as it fluctuates randomly, explores the configuration space defined by the 16 states. It periodically gets "stuck" in particular configurations or evolves between certain arrangements, and its speed of fluctuation will depend on the temperature we have set in the simulation. But symmetry-breaking behaviour does not become really apparent until we increase the system size.

Practically, analysing the Ising model becomes increasingly difficult as N increases. When $N = 2$, we have 16 configurations. In general, the number of configurations is 2^{N^2}, and this grows very fast with N. Thus, when $N = 10$, the number of configurations is around 10^{30}, well beyond the reach of humanity's computing resources to store all of them (though trivial to write down just one of those configurations; see Figure 19.4). New methods are needed to gain meaningful insight into this sort of problem and, as explained later, the renormalization group is one such method.

Phase transitions exhibit an infinitely sharp discontinuity as a function of temperature. This arises from the supposed infinite size of the system. In finite systems the partition function $Z = \Sigma \exp(-E_i/k_BT)$ (on which an analysis of $\langle M(T)\rangle$ can be based) exhibits no singularities, and we expect only a smooth crossover between an aligned state at zero temperature and a disordered one at infinite temperature. Phase transitions in nature are, of course, exhibited in finite systems, so can they be said to be really phase transitions? Strong arguments can be given to answer "yes" to this question. A typical macroscopic lump of matter contains a large number N of atoms, perhaps around 10^{23}, and so the discontinuity in properties will occur over such a narrow temperature

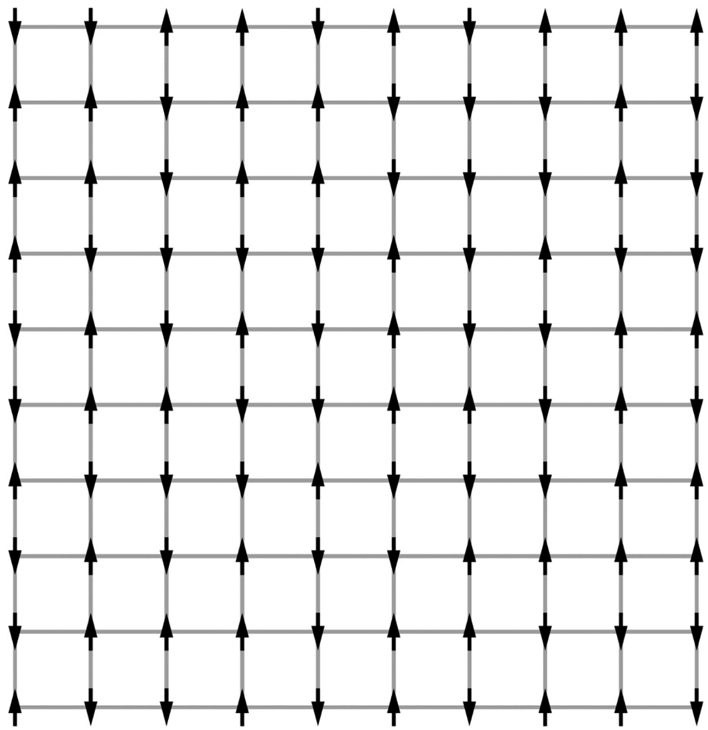

Figure 19.4 One of the approximately 10^{30} configurations for a 10×10 grid of Ising spins

interval that its width is effectively impossible to measure. Moreover, when simulations of finite systems are performed, it is possible to study how the sharpness of the phase transition scales with system size. Such studies show that the temperature range of the crossover falls dramatically as N increases, and so we can describe the phenomenon as a limiting process as N tends towards infinity. The mathematics is most straightforward in the infinite limit, but the finite-size corrections to this model are usually negligible for practical purposes (except in very small systems), but more importantly, they are well understood.

However, we don't need the infinite N limit to be fully made to discern the extraordinary properties of broken symmetry states. In a typical macroscopic lump of matter $N \approx 10^{23}$, and that is big enough. A key insight (Anderson 1984) is the realization that once in a broken symmetry state, it is very difficult to reach other broken symmetry states via fluctuations when the number of particles becomes large. In fact, the probability of all N particles fluctuating at once is proportional to $\exp(-\sqrt{N})$ and becomes negligible as N becomes large. The broken symmetry state is not an eigenstate of the system Hamiltonian H (it breaks one of the symmetries of H) and so it is not a stationary state. Yet it is a stable state of the system (stable, at least, on timescales greater than the age of the universe), an emergent feature of the large-number nature of matter. How do these states form? In an infinite system, an infinitesimal perturbation is sufficient to cause the system to choose one of the symmetry-broken states. Such states can then be modelled by introducing small symmetry-breaking fields (to break the symmetry), and then taking the limit such that N tends to infinity before the symmetry-breaking field is removed; taking the limits in the reverse order leads to a situation in which fluctuations diverge (van Wezel and van den Brink 2006).

Consequences for symmetry breaking

When symmetry breaking occurs, we will expect a sharp change of behaviour at the critical temperature. The region near the phase transition is called the *critical region*. Having broken the symmetry, the system will have a strong energetic preference for staying in that particular broken symmetry state, and attempts by the experimenter to change the choice the system has made in breaking the symmetry will meet with resistance. Thus crystals (which have broken translation symmetry) don't bend or stretch very easily. Ferromagnets show permanent magnetism, and the spins are said to be "stiff". These *rigidity phenomena* are common to broken symmetry states. If symmetry is broken differently in two adjacent parts of a macroscopic sample, the boundary will contain a defect (e.g. a dislocation in a crystal or a domain wall in a ferromagnet).

There are further consequences for the excitations of these symmetry-broken phases. At $T = 0$, the system is perfectly ordered. At finite temperature this order is weakened by excitations in the order parameter. In crystals these excitations are called lattice waves, while in ferromagnets the analogous modes are called spin waves. These wave-like phenomena are normal modes of the system and are therefore susceptible to similar analysis to that applied to the harmonic oscillator, so that quantum mechanically these modes can only acquire energy in discrete lumps. Thus, the energy in these wave systems behaves like a collection of particles. The quanta (i.e. the quantized particles) of lattice waves are termed phonons; those in spin waves are termed magnons. To this menagerie, we can add rotons (quantized rotational waves in liquid helium), ripplons (quantized surface waves), plasmons (quantized plasma waves) and polaritons (particles resulting from coupled light and lattice waves). The mathematical structure underlying these new particles forming in condensed matter systems is entirely analogous to that used to describe the quantization of the electromagnetic waves that results in photons.

So now we come to the key question: Are these emergent particles real? From the perspective of quantum field theory, the answer is a resounding yes. Each of these particles emerges from a wave-like description in a manner that is entirely analogous to that of photons. These emergent particles behave like particles: you can scatter other particles off them. Electrons will scatter off phonons, an interaction that is involved in superconductivity. Neutrons can be used to study the dispersion relation of both phonons and magnons using inelastic scattering techniques. Yes, they are the result of a collective excitation of an underlying substrate. But so are "ordinary" electrons and photons, which are excitations of quantum field modes.

Thus, all of these excitations that arise as a consequence of broken symmetry possess the features you expect of an emergent quantity: they are made up of the combination of another apparently different type of thing but have the properties of something new. Thus, for example, the phonon has a particle-like nature that can propagate through a crystal, but it is entirely composed of atoms. In the same way, a Mexican wave in a crowd has a life of its own, even though it is entirely composed of individuals with their hands in the air, none of whom could do a Mexican wave on his or her own. These emergent particles fall into the class of what one can describe as "real patterns" (Dennett 1991), thereby deserving their metaphysical status from the intellectual coherence of their description of physical phenomena. Just as the "life-forms" in Conway's Game of Life are not simply mirages imagined in a sea of flickering pixels (see Blundell 2017), so these condensed matter particles which emerge from broken symmetry states can be legitimately claimed to be *real* and not just *apparent*. Just as real, in fact, as the electrons that emerge from the vacuum as excitations in the electron quantum field.

The renormalization group and universality

The renormalization group breaks big problems down into small ones. This is easy to understand in a magnetic system of $N \times N = N^2$ spins by considering the Kadanoff block-spin transformation, in which na $m \times m$ cluster of m^2 sites is considered together and the majority spin in the cluster measured (Kadanoff 1966). The old $N \times N$ lattice is then replaced by a new $(N/m) \times (N/m)$ lattice that is now composed of these "block spins" that each represent the average behaviour of the cluster, thinking of it as a unit in itself. This transformation has therefore acted like a blurring process, reducing the information in the system and only preserving the large-scale structure. One might guess that the couplings in this reduced system are identical to the original system, but usually this is not the case. Kadanoff's block-spin transformation can then be applied repeatedly in a series of successive blurrings, and one can see how the couplings develop as a consequence.

This is the key idea behind the so-called renormalization group procedure, which is to look at how a model behaves when the scale changes (and naturally generalizes Kadanoff's original argument). Typically, the rescaling procedure results in identical physics but with altered coupling constants. Then with repeated rescaling we can study how the coupling constants change (or to use the lingo of the subject: how they "flow"). The coupling constants g_i describe a point (g_1, $g_2, g_3, \ldots$) in a multidimensional space, and repeated rescalings result in a trajectory through this space, known as the renormalization group flow. In some cases, the rescalings cause the point to be blasted out to infinity, in others to converge on to a fixed point. At a fixed point, the coupling constants are invariant to further rescalings and describe scale-invariant physics.

The renormalization group is not actually a group in the mathematical sense. The transformations consisting of rescaling and integrating up to a cut-off have no inverse because the integration process results in blurring and a loss of detail. The renormalization group should probably be called the renormalization semigroup, but sometimes, contradictory terminology sticks (as Groucho Marx quipped, think of "military intelligence"). Nevertheless, the renormalization group procedure provides important insights because it shows quantitatively how fine-scale structure is progressively ignored and the physics of critical phenomena depend on these larger-scale, what you might call "structural", features of the theory. It suddenly becomes clear why certain measurable features of the physics of many systems are seemingly independent of the microscopic details, why phase transitions in certain magnets can share common features with the liquid–solid transition or with transitions that can occur in helium films or in traffic flow.

These common features, which include scaling relations of the order parameter and correlation lengths, give rise to the concept of "universality". This is the notion that some of the key parameters of a phase transition depend only on some broad characterizations of the system but not on certain fiddly details. Sometimes those fiddly details are not so obviously ignorable. For example, it might be important to know whether you are dealing with a two-dimensional lattice or a three-dimensional one, but not whether that two-dimensional lattice contains points on a square or a triangular grid. Whether you are studying a square or a triangular lattice is a detail that is important on the scale of the lattice, but its importance diminishes as your perspective widens and you examine the properties of the system on a larger and larger scale. However, whether the lattice itself is two- or three-dimensional is important on all scales.

The features that determine which "university class" will apply to the problem in hand are (i) the dimensionality of its order parameter field, (ii) the dimensionality of the system and (iii) simply whether the interactions are short or long range. They do not depend on the details of the microscopic interactions between constituent parts or whether the system is made up of polymers in a beaker or particles in the early universe. Thus, for example, the properties of a fluid near its critical point and a three-dimensional Ising ferromagnet are both in the same universality class. Both have one-dimensional order parameter fields (the difference between the liquid and vapour densities, $\rho_v - \rho_l$, in the case of the fluid; the magnetization M, in the case of the Ising ferromagnet) that are defined in a three-dimensional space, and both vanish as the temperature is increased above the critical temperature. Both have a corresponding response function (the compressibility κ or the susceptibility χ, respectively) that diverge in the same way (i.e. with the same critical exponents) as the critical point is reached. The "multiple realizability" of phenomena produced at critical points (Batterman 2001) is a feature of the oft-neglected role of singularities in physics and the importance of considering the asymptotic behaviour close to the critical point (Berry 1994), exactly what the renormalization group aims to do. Here I want to underline the way in which critical behaviour *emerges* from a substrate of microscopic interactions and that the predominant narrative for understanding such behaviour operates almost exclusively at that higher, emergent level. In many physical situations, the emergent description simplifies a more complex lower-level account; here, the emergent description makes the lower-level account irrelevant.

Conclusion

Broken symmetry phases are all around us and include crystals, magnets, liquid crystals and superconductors. By breaking symmetry, these phases forfeit the status of being a "stationary state" of the sort beloved of elementary quantum mechanics treatments. Nevertheless, because they are composed of a macroscopic number of particles, these phases can be stable on timescales longer than the universe and give rise to a range of emergent ordered properties. The excitations of such phases give rise to a new set of emergent particles, such as phonons and magnons. Theoretical approaches, such as the renormalization group method, allow one to study in detail how certain microscopic parameters become irrelevant in the critical regime, while other features remain crucially relevant. Broken symmetry can therefore be seen as a well-studied paradigm of emergent behaviour in the physical world.

References

Altland, A. and Simons, B. (2006) *Condensed Matter Field Theory*, Cambridge: Cambridge University Press.

Anderson, P. W. (1984) *Basic Notions of Condensed Matter Physics*, Menlo Park: Benjamin-Cummings.

Batterman, R. W. (2001) *The Devil in the Details: Asymptotic Reasoning in Explanation, Reduction and Emergence*, Oxford: Oxford University Press.

Berry, M. (1994) "Asymptotics, singularities and the reduction of theories", in *Proc. 9th Int. Cong. Logic, Method., and Phil. of Sci.* 9: 597–607, eds. D. Prawitz, B. Skyrms, and D. Westerståhl, Amsterdam: Elsevier.

Binney, J. J., Dowrick, N. J., Fisher, A. J., and Newman, M. E. J. (1992) *The Theory of Critical Phenomena*, Oxford: Oxford University Press.

Blundell, S. J. (2001) *Magnetism in Condensed Matter* (2nd edition), Oxford: Oxford University Press.

Blundell, S. J. (2017) "Emergence, causation and storytelling: Condensed matter physics and the limitations of the human mind", *Philosophica* 92: 139–164.

Blundell, S. J. and Blundell, K. M. (2010) *Concepts in Thermal Physics*, Oxford: Oxford University Press.

Delamotte, B. (2004) "A hint of renormalization", *Am. J. Phys.* 72: 170–184.

Dennett, D. C. (1991) "Real patterns", *J. Philosophy* 88: 27–51.

Fisher, M. E. (1974) "The renormalization group in the theory of critical behavior", *Rev. Mod. Phys.* 46: 597–616.

Goldenfeld, N. (1992) *Lectures on Phase Transitions and the Renormalization Group*, Boulder, CO: Westview Press.

Kadanoff, L. P. (1966) "Scaling laws for Ising models near T_c", *Physics* 2: 263–272.

Kadanoff, L. P. (2013) "Theories of matter: Infinities and renormalization", in R. Batterman (ed.) *The Oxford Handbook of Physics and Philosophy*, Oxford: Oxford University Press, pp. 141–188.

Lancaster, T. and Blundell, S. J. (2014) *Quantum Field Theory for the Gifted Amateur*, Oxford: Oxford University Press.
Le Bellac, M., Mortessagne, F., and Batrouni, G. G. (2004) *Equilibrium and Non-Equilibrium Statistical Thermodynamics*, Cambridge: Cambridge University Press.
Ma, S.-K. (1973) "Introduction to the renormalization group", *Rev. Mod. Phys.* 45: 589–614.
McComb, W. D. (2004) *Renormalization methods*, Oxford: Oxford University Press.
Nishimori, H. and Ortiz, G. (2011) *Elements of Phase Transitions and Critical Phenomena*, Oxford: Oxford University Press.
Stinchcombe, R. B. (1989) "Fractals, phase transitions and criticality", *Proc. Roy. Soc. London* 423: 17–33.
Tolédano, J.-C. and Tolédano, P. (1987) *The Landau Theory of Phase Transitions*, Singapore: World Scientific.
van Wezel, J. and van den Brink, J. (2006) "Spontaneous symmetry breaking in quantum mechanics", *Am. J. Phys.* 75: 635–638.
Weinberg, S. (1995) *The Quantum Theory of Fields*, Cambridge: Cambridge University Press.
Wilson, K. G. (1972) "Feynman-graph expansion for critical exponents", *Phys. Rev. Lett.* 28: 548–551.
Wilson, K. G. (1983) "The renormalization group and critical phenomena", *Rev. Mod. Phys.* 78: 583–600.
Wilson, K. G. and Fisher, M. E. (1972) "Critical Exponents in 3.99 Dimensions", *Phys. Rev. Lett.* 28: 240–243.
Wilson, K. G. and Kogut, J. (1973) "The renormalization group and the ε expansion", *Phys. Rep.* 12(2): 75–200.
Zinn-Justin, J. (2007) *Phase Transitions and Renormalization Group*, Oxford: Oxford University Press.

Further reading

A very good introduction to the renormalization group may be found in Nishimori and Ortiz (2011) and McComb (2004). In the context of quantum field theory, a non-technical account may be found in Lancaster and Blundell (2014) and Altland and Simons (2006), with more details in Weinberg (1995). Good reviews of the renormalization group in critical phenomena may be found in chapter 4 of Le Bellac et al. (2004), and also in Binney et al. (1992), Ma (1973), Fisher (1974), Wilson (1983), Wilson and Kogut (1973) and Zinn-Justin (2007). A particularly clear exposition is given in Goldenfeld (1992), and historical background is covered in Kadanoff (2013). Philosophical implications of asymptotic limits are described in Batterman (2011), and Berry (1994) provides the scientific background. The implications of broken symmetry to condensed matter systems are described in Anderson (1984), though see Blundell (2001) for applications to magnetism. For a technical account of the relationship between symmetry groups and phase transitions, see Tolédano and Tolédano (1987).