1 · Preview

1.1 Apologia

The lectures on which this book is based were intended for a 'mixed audience'. According to the context, that phrase might have certain social connotations, but here it implies a more fundamental distinction: some of the audience were basically physicists, and others were basically mathematicians. This distinction, between those who think in terms of real objects and those who deal in abstract ideas, is an unfortunate fact of scientific life today.

The desire to be intelligible to two classes of student has been my main preoccupation in preparing the lectures and writing the book. Consequently, any reader will probably find some material which (to him) is tiresome and elementary; such material is included for the benefit of other readers, in the cause of scientific harmony. I have tried to prescribe a proper dose of generality - not too much to discourage those who have a particular application in mind, nor too little for those who wish to see the underlying structure.

The book has five chapters, each subdivided into sections. The first chapter is intended as a broad introduction to the subject, and it is written in a more informal manner than the rest. There are two short appendices at the end of the book, and these are referred to in Chapters 2, 3 and 4. Apart from this, there are no references in the main text; notes and references for each chapter are given at the end of the chapter.

1.2 States on a graph

Our main object of study is a finite system of particles, some of which interact in pairs. We shall not be concerned with the nature of the particles, their relative positions, or the kind of interactions; such

things are important when one studies a particular physical structure, but they do not lie at the heart of the matter. In order to describe the pattern of interactions among the particles of the system we shall require a little of the terminology of graph theory.

A **graph** consists of a set of **vertices** and a set of **edges**, together with an **incidence relation**: each edge is **incident** with either one or two vertices. Graphs are often represented by diagrams in which the vertices are points and the edges are line segments joining the relevant points (see Fig. 1). An edge incident with just one vertex is called a **loop**, and is represented as such. Sets of edges incident with the same pair of vertices are called **multiple edges**. A graph with no loops or multiple edges is said to be **simple** (Fig. 1). We shall adopt the convention that the term 'graph' will always mean 'simple graph', unless the context indicates otherwise.

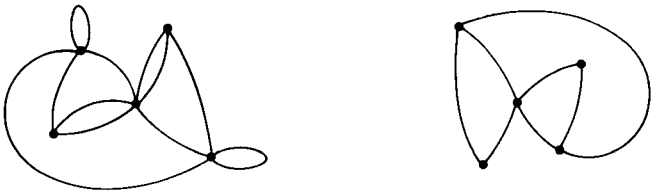

Fig. 1. (a) A graph (b) A simple graph

In the ensuing theory it will be usual to think of the vertices of a graph as a set of particles, where the edges signify the existence of interactions between some pairs of particles.

We shall reserve the notations V and E for the vertex-set and edge-set of a graph; occasionally, when the graph G needs to be named explicitly, VG and EG will be used. A graph is said to be **connected** if any pair v, w of its vertices may be linked by a **path**, that is, a sequence

$$v = v_0,\ e_1,\ v_1,\ e_2,\ \ldots,\ e_r,\ v_r = w$$

of alternate vertices and edges, where e_i is an edge incident with v_{i-1} and v_i $(1 \leq i \leq r)$. If a graph is not connected, it falls into a number of components, each of which is connected. In terms of a system of interacting particles, we may separate the system into subsystems, each of which is independent of the rest and cannot be further subdivided.

Let A be a finite set, and G a graph. A function ω from the vertex-set V of G to the set A (written $\omega : V \to A$) assigns to each vertex an element of A; that is, each particle is given some 'attribute' or 'configuration' or 'colour'. A function ω thus defines a state of the system of particles represented by the graph, and the set of all such states will be denoted by $\Omega(G, A)$, or just Ω.

For theoretical reasons it is convenient to endow the set A with some algebraic structure. The most apt structure is that of a 'ring'; that is, we postulate the existence of two operations, $+$ and $\cdot$, satisfying the usual rules of arithmetic, except that division (the inverse of the $\cdot$ operation) is not allowed. All the rings that we consider will contain elements 0 and 1, with the usual properties of those symbols, and both operations will be commutative. For each positive integer m, the set of residues modulo m forms a ring, with the usual operations of modular arithmetic. We shall denote this special ring by A_m. Thus any non-empty finite set is in one-to-one correspondence with at least one ring. It follows that a ring structure may be imposed on any finite set of attributes, without loss of generality, and with some gain. For many purposes (but not all) the rings of residues are a sufficient level of generality.

When A is a ring the states ω in $\Omega(G, A)$ may themselves be combined by operations derived from the structure of A. In fact, if ω_1 and ω_2 are such states, then we may define states $\omega_1 + \omega_2$ and $\omega_1 . \omega_2$ by the rules

$$(\omega_1 + \omega_2)(v) = \omega_1(v) + \omega_2(v) ,$$
$$(\omega_1 . \omega_2)(v) = \omega_1(v) . \omega_2(v) .$$

The set Ω thus becomes a ring itself. In particular, there is a 0-state 0 defined by $0(v) = 0$ for all v in V, where the second 0 is the zero

element of A, and similarly a 1-state defined by $1(v) = 1$.

In a manner analogous to that used to define Ω, we may introduce the ring $\Phi = \Phi(G, A)$ of functions $\phi : E \to A$. We think of a function ϕ in Φ as an assignment of a 'flow' $\phi(e)$ to each edge e of G. In order to describe the 'direction' of such a flow, we must choose an **orientation** of the graph G; that is, for each edge e of G one of the two incident vertices is chosen to be the 'positive end' of e, and the other is chosen to be the 'negative end'. We make the convention that the single vertex incident with a loop is its positive end. An orientation is usually represented on a diagram by placing an arrow on each edge, pointing towards its positive end (Fig. 2). Although the introduction of

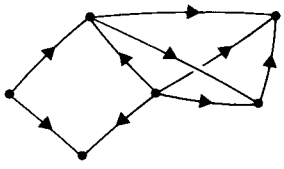

Fig. 2. A graph with an orientation

an orientation is necessary in order to yield satisfactory definitions, the actual orientation chosen is immaterial.

Given a finite graph G with an orientation, we define a matrix D with $|V|$ rows, labelled by the vertices of G, and $|E|$ columns, labelled by the edges of G, as follows:

$$D_{ve} = \begin{cases} 1 & \text{if } v \text{ is the positive end of } e; \\ -1 & \text{if } v \text{ is the negative end of } e; \\ 0 & \text{if } v \text{ is not incident with } e. \end{cases}$$

D is called the **incidence matrix** of G, with respect to the chosen orientation. The entries of D are taken to be elements of the ring A; every ring has elements 0, 1, and -1, although in the case of the residues modulo 2 (for example) 1 and -1 are the same.

We can now define two important operators. They are:

the **boundary** operator $\partial : \Phi(G, A) \to \Omega(G, A)$

and the **coboundary** operator $\delta : \Omega(G, A) \to \Phi(G, A)$.

The operator ∂ assigns to each flow ϕ on G a state $\partial\phi$, and δ assigns to each state ω a flow $\delta\omega$, defined by

$$(\partial\phi)(v) = \sum_{e \in E} D_{ve} . \phi(e); \qquad (\delta\omega)(e) = \sum_{v \in V} D_{ve} . \omega(v).$$

Intuitively, the value of $(\partial\phi)(v)$ represents the net accumulation of flow at the vertex v. As for $\delta\omega$, if e is not a loop then there are just two non-zero terms D_{ve} : a term D_{xe} corresponding to the positive end x of e ($D_{xe} = 1$), and a term $D_{ye} = -1$ corresponding to the negative end y of e. Consequently

$$(\delta\omega)(e) = \omega(x) - \omega(y),$$

and $(\delta\omega)(e)$ represents the difference in the values of ω at the ends of e.

We shall introduce other notions from graph theory as they are needed.

1.3 Interaction models

In this section we shall set up a mathematical structure, called an interaction model, which is the main topic in the remainder of this book. The physical background motivating the formulation of the model will be surveyed in the next section.

For the time being, all graphs considered will be simple and finite. We begin with the set $\Omega = \Omega(G, A)$ of states on a graph G with values in the ring A. To each state $\omega : V \to A$ we assign a complex number $I(\omega)$ called the weight of ω; in other words I is a function

$$I : \Omega(G, A) \to \mathbf{C},$$

where $\mathbf{C}$ denotes the complex numbers. In many instances the values of I will be real and non-negative. When that is the case we may put

$$Z = \sum_{\omega \in \Omega} I(\omega)$$

and interpret the quantity $I(\omega)/Z$ as being the probability that the particles of the system have the attributes specified by ω.

We shall be concerned with weights which are derived from the local structure of the graph G in the following way. Any state ω has a coboundary $\delta\omega$; as explained in the previous section, the value $\delta\omega(e)$ represents the difference in the values of ω at the ends of the edge e. We introduce an 'interaction function' $i : A \to C$, whose value $i(a)$ is taken to represent the strength of the interaction between two particles when the difference (in the ring A) of their 'attributes' is a. Because the interaction should not depend on the order in which the particles are considered, we shall always assume that i satisfies the symmetry condition $i(a) = i(-a)$. When the state of the system is specified by ω, the interaction on the edge e of G is given by $i[\delta\omega(e)]$; because of the symmetry condition, this does not depend on the orientation used in defining δ. We shall take the weight of ω to be the product of these terms over all edges of G. (The reason for the occurrence of a product, rather than a sum, will appear shortly.) We are now in a position to recast the foregoing discussion into a set of basic definitions.

An **interaction model** $\mathfrak{M}$ consists of a ring A and an **interaction function** $i : A \to C$, with the property $i(a) = i(-a)$. We write $\mathfrak{M} = (A, i)$, signifying that $\mathfrak{M}$ is an ordered pair.

If we are given an interaction model $\mathfrak{M}$ and a graph G, we shall speak of the interaction model $\mathfrak{M}$ on G, and associate with $\mathfrak{M}$ and G the **weight function**

$$I(\omega) = \prod_{e \in E} i[\delta\omega(e)] , \tag{1.3.1}$$

and the **partition function**

$$Z(\mathfrak{M}, G) = \sum_{\omega \in \Omega} I(\omega) = \sum_{\omega \in \Omega} \prod_{e \in E} i[\delta\omega(e)]. \tag{1.3.2}$$

For example, let us consider the very simple model $\mathfrak{C}$ whose interaction function is defined by

$$i(a) = \begin{cases} 1 & \text{if } a \neq 0, \\ 0 & \text{if } a = 0. \end{cases}$$

For this function, $i[\delta\omega(e)] = 0$ when ω assigns the same element of A to the two ends of e, and it is equal to 1 otherwise. Consequently $I(\omega)$ is zero unless ω assigns different 'colours' to the two ends of each edge in G, when it is equal to 1. If we use the above notation, then $Z(\mathbb{C}, G)$ is the number of 'proper colourings' of G with $|A|$ colours, where the adjective 'proper' signifies that adjacent vertices have different colours. This function has been much studied in graph theory, under the name of the 'chromatic polynomial'.

In physics, the models studied usually have the property that the interaction $i(a)$ is real and positive, for all a in A. We shall say that an interaction model is **positive** if its interaction function has this property. Such interaction functions may be written as

$$i(a) = \exp j(a) \qquad (a \in A),$$

for some real-valued function j defined on A. The corresponding weight function I is then of the form

$$I(\omega) = \exp J(\omega) \qquad (\omega \in \Omega),$$

where

$$J(\omega) = \sum_{e \in E} j[\delta\omega(e)] .$$

These remarks (together with the physical principles described in the next section) explain how the sum of interactions becomes a product in our general formulation.

Interactions of the kind described in the preceding paragraph have an interesting additional property, related to the probability interpretation of the weight function. Suppose that I is any real, positive weight function on Ω, and extend I to the subsets of Ω by putting

$$I(\Lambda) = \sum_{\omega \in \Lambda} I(\omega) \qquad (\emptyset \neq \Lambda \subseteq \Omega),$$

$$I(\emptyset) = 0.$$

Then I/Z is a probability measure on the subsets of Ω, and we may

define conditional probabilities in the usual way. That is, the conditional probability that a state belongs to a subset X of Ω, given that it belongs to Y, is $I(X \cap Y)/I(Y)$.

Let v be a given vertex of a simple graph G. We shall say that a real positive weight function I has the **Markov property** if the following holds: the conditional probability that a state is ω, given that it agrees with ω on all vertices of G except v, depends only on the values of ω at v and the vertices adjacent to v in G. Roughly speaking, the probability that the system has a particular configuration at v depends only on its values at the neighbours of v.

It is fairly easy to see that any weight function I, which arises from a positive interaction model $\mathcal{P} = (A, i)$, has the Markov property. In view of the fact that interactions occur only between neighbouring vertices, this is perhaps not too surprising, but a proof seems called for. First, it is clear from (1.3.1) that I is real and positive whenever i is. Now the conditional probability occurring in the definition of the Markov property is just

$$I(\omega)/\Sigma I(\theta),$$

where the sum is over all those states θ which agree with ω except at v. These are the states ω_a $(a \in A)$ defined by

$$\omega_a(v) = \omega(v) + a, \quad \omega_a(w) = \omega(w) \qquad (w \neq v).$$

In particular, $\omega = \omega_0$. Now

$$\begin{aligned} I(\omega_a) &= \prod_{e \in E} i[\delta\omega_a(e)] \\ &= (\text{product over edges } e \text{ incident with } v) \\ &\quad \times (\text{product over } e \text{ not incident with } v). \end{aligned}$$

Since ω_a is independent of a except at v, $\delta\omega_a$ is independent of a, except on those edges incident with v. Consequently, in the conditional probability

$$I(\omega)/\sum_{a \in A} I(\omega_a)$$

the product terms involving edges not incident with v may be cancelled throughout. The remaining terms depend only on the values of ω at v and its neighbours, and so I has the Markov property.

A partial converse of the preceding result is also true. If G satisfies some simple conditions, then given any weight function I on $\Omega(G, A)$ which has the Markov property, there is some interaction function i such that I is derived from i by the product formula (1.3.1). This converse enables us to use the interaction model in the study of stochastic processes on graphs. The equilibrium distributions of such processes often turn out to have the Markov property; consequently, they may be described in terms of our interaction model.

1.4 Physical background

An interaction model whose weight function is of the form $I(\omega) = \exp J(\omega)$ is typical of the models studied in statistical mechanics. The usual distribution of states is the Gibbs canonical distribution, wherein the weight of a state ω is equal to

$$\exp\{(-1/kT)H(\omega)\}. \qquad (1.4.1)$$

Here k is an absolute constant, T is the temperature, and $H(\omega)$ is a Hamiltonian function representing the energy of the state ω. In general, $H(\omega)$ will be a sum of terms corresponding to individual interactions, and so our product formula for I is obtained, as mentioned in the previous section. The various thermodynamic quantities, such as free energy and specific heat, may be derived from the partition function considered as a function of T.

For the purposes of illustration we shall discuss a famous example of an interaction model - the Ising model of ferromagnetism. In physical terms, the particles are the atoms of a ferromagnetic substance, each of which has a magnetic moment or 'spin'. There are just two possible configurations for each spin, and they are conventionally thought of as 'up' and 'down'. Like spins contribute an amount of energy $-L$ to the Hamiltonian, and unlike spins contribute $+L$.

In our notation, the two configurations 'up' and 'down' become the two elements 0 and 1 of the ring A_2 of residues modulo 2. In accordance with the Gibbs distribution, the weight function is of the form

$$I(\omega) = \exp\{(-1/kT)\Sigma(\pm L)\}\ .$$

So if we put $\varepsilon = \exp(L/kT)$ and define an interaction function i on A_2 by the rules

$$i(0) = \varepsilon, \qquad i(1) = \varepsilon^{-1},$$

we have precisely the weight function of the **Ising model** $\mathcal{I} = (A_2, i)$, in terms of our definition (1.3.1). In fact, we have a set of models, one for each positive value of the temperature T.

It is fairly easy to find explicit expressions for the partition function of the Ising model on various well-known graphs. For example, consider the **complete graph** K_n, which has n vertices and $\frac{1}{2}n(n-1)$ edges, one edge joining each pair of distinct vertices; this graph corresponds to a physical system in which each pair of particles interacts. If the state ω on K_n has l up-spins and $n-l$ down-spins, then there are $l(n-l)$ unlike pairs and $\frac{1}{2}n(n-1) - l(n-l)$ like pairs; so the weight of ω is

$$I(\omega) = \varepsilon^{\frac{1}{2}n(n-1)-2l(n-l)}.$$

The partition function is

$$Z(\mathcal{I}, K_n) = \varepsilon^{\frac{1}{2}n(n-1)} \sum_{l=0}^{n} \binom{n}{l} \varepsilon^{-2l(n-l)}. \qquad (1.4.2)$$

There are a couple of fairly simple reduction formulae which can be employed to advantage. First, if G is a disconnected graph with (say) two components G_1 and G_2 then for any model $\mathfrak{M}$ we have

$$Z(\mathfrak{M}, G) = Z(\mathfrak{M}, G_1)Z(\mathfrak{M}, G_2). \qquad (1.4.3)$$

This may be proved directly from the definitions, just by noting that each state on G splits into a state on G_1 and a state on G_2, and there are no