10. The spatial variation of this field was carefully analyzed† and found to be entirely consistent with classical theory out to a distance of at least 10^5 kilometers (km) from the planet. This is tantamount to a test, albeit indirect, of Coulomb's law over that distance.

To summarize, we have every reason for confidence in Coulomb's law over the stupendous range of 24 decades in distance, from 10^{-14} to 10^{10} cm, if not farther, and we take it as the foundation of our description of electromagnetism.

ENERGY OF A SYSTEM OF CHARGES

1.5 In principle, Coulomb's law is all there is to electrostatics. Given the charges and their locations we can find all the electrical forces. Or given that the charges are free to move under the influence of other kinds of forces as well, we can find the equilibrium arrangement in which the charge distribution will remain stationary. In the same sense, Newton's laws of motion are all there is to mechanics. But in both mechanics and electromagnetism we gain power and insight by introducing other concepts, most notably that of energy.

Energy is a useful concept here because electrical forces are *conservative*. When you push charges around in electric fields, no energy is irrecoverably lost. Everything is perfectly reversible. Consider first the work which must be done *on* the system to bring some charged bodies into a particular arrangement. Let us start with two charged bodies or particles very far apart from one another, as indicated at the top of Fig. 1.4, carrying charges q_1 and q_2. Whatever energy may have been needed to create these two concentrations of charge originally we shall leave entirely out of account. Bring the particles slowly together until the distance between them is r_{12}. How much work does this take?

It makes no difference whether we bring q_1 toward q_2 or the other way around. In either case the work done is the integral of the product: force times displacement in direction of force. The force that has to be applied to move one charge toward the other is equal to and opposite the Coulomb force.

$$W = \int \text{force} \times \text{distance} = \int_{r=\infty}^{r_{12}} \frac{q_1 q_2(-dr)}{r^2} = \frac{q_1 q_2}{r_{12}} \qquad (3)$$

Because r is changing from ∞ to r_{12}, the increment of displacement is $-dr$. We know the work done on the system must be positive for charges of like sign; they have to be pushed together. With q_1 and q_2 in esu, and r_{12} in cm, Eq. 3 gives the work in ergs.

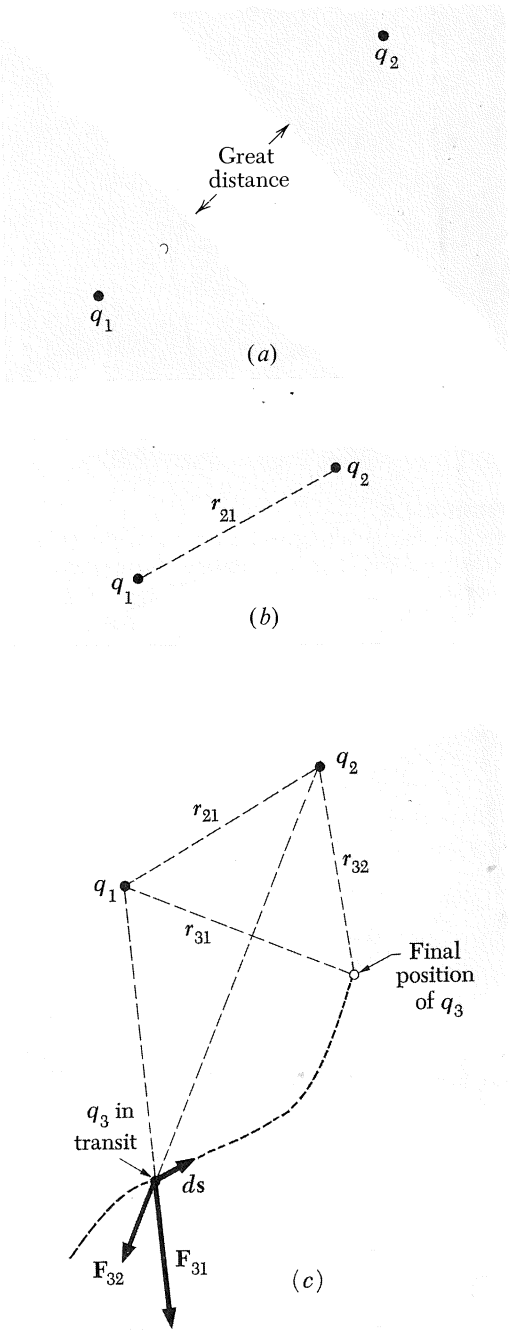

FIGURE 1.4
Three charges are brought near one another. First q_2 is brought in; then with q_1 and q_2 fixed, q_3 is brought in.

†L. Davis, Jr., A. S. Goldhaber, M. M. Nieto, *Phys. Rev. Lett.* **35**:1402 (1975). For a review of the history of the exploration of the outer limit of classical electromagnetism, see A. S. Goldhaber and M. M. Nieto, *Rev. Mod. Phys.* **43**:277 (1971).

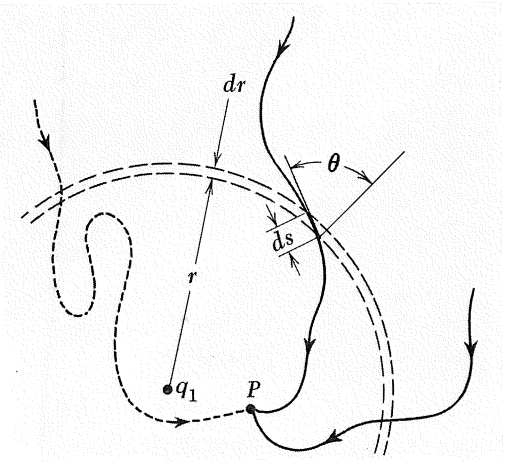

FIGURE 1.5
Because the force is central, the sections of different paths between $r + dr$ and r require the same amount of work.

This work is the same whatever the path of approach. Let's review the argument as it applies to the two charges q_1 and q_2 in Fig. 1.5. There we have kept q_1 fixed, and we show q_2 moved to the same final position along two different paths. Every spherical shell such as the one indicated between r and $r + dr$ must be crossed by both paths. The increment of work involved, $-\mathbf{F} \cdot d\mathbf{s}$ in this bit of path, is the same for the two paths.† The reason is that $\mathbf{F}$ has the same magnitude at both places and is directed radially from q_1, while $d\mathbf{s} = dr/\cos\theta$; hence $\mathbf{F} \cdot d\mathbf{s} = F\,dr$. Each increment of work along one path is matched by a corresponding increment on the other, so the sums must be equal. Our conclusion holds even for paths that loop in and out, like the dotted path in Fig. 1.5. (Why?)

Returning now to the two charges as we left them in Fig. 1.4*b*, let us bring in from some remote place a third charge q_3 and move it to a point P_3 whose distance from charge 1 is r_{31} cm, and from charge 2, r_{32} cm. The work required to effect this will be

$$W_3 = -\int_{\infty}^{P_3} \mathbf{F}_3 \cdot d\mathbf{s} \tag{4}$$

Thanks to the additivity of electrical interactions, which we have already emphasized,

$$\begin{aligned} -\int \mathbf{F}_3 \cdot d\mathbf{s} &= -\int (\mathbf{F}_{31} + \mathbf{F}_{32}) \cdot d\mathbf{s} \\ &= -\int \mathbf{F}_{31} \cdot d\mathbf{r} - \int \mathbf{F}_{32} \cdot d\mathbf{r} \end{aligned} \tag{5}$$

That is, the work required to bring q_3 to P_3 is the sum of the work needed when q_1 is present alone and that needed when q_2 is present alone.

$$W_3 = \frac{q_1 q_3}{r_{31}} + \frac{q_2 q_3}{r_{32}} \tag{6}$$

The total work done in assembling this arrangement of three charges, which we shall call U, is therefore

$$U = \frac{q_1 q_2}{r_{12}} + \frac{q_1 q_3}{r_{13}} + \frac{q_2 q_3}{r_{23}} \tag{7}$$

We note that q_1, q_2, and q_3 appear symmetrically in the expression above, in spite of the fact that q_3 was brought up last. We would have reached the same result if q_3 had been brought in first. (Try it.) Thus U is independent of the *order* in which the charges were assem-

†Here we use for the first time the scalar product, or "dot product," of two vectors. A reminder: the scalar product of two vectors $\mathbf{A}$ and $\mathbf{B}$, written $\mathbf{A} \cdot \mathbf{B}$, is the number $AB\cos\theta$. A and B are the magnitudes of the vectors $\mathbf{A}$ and $\mathbf{B}$, and θ is the angle between them. Expressed in terms of cartesian components of the two vectors, $\mathbf{A} \cdot \mathbf{B} = A_xB_x + A_yB_y + A_zB_z$.

bled. Since it is independent also of the route by which each charge was brought in, U must be a unique property of the final arrangement of charges. We may call it the *electrical potential energy* of this particular system. There is a certain arbitrariness, as always, in the definition of a potential energy. In this case we have chosen the zero of potential energy to correspond to the situation with the three charges already in existence but infinitely far apart from one another. The potential energy *belongs to the configuration as a whole.* There is no meaningful way of assigning a certain fraction of it to one of the charges.

It is obvious how this very simple result can be generalized to apply to any number of charges. If we have N different charges, in any arrangement in space, the potential energy of the system is calculated by summing over all pairs, just as in Eq. 7. The zero of potential energy, as in that case, corresponds to all charges far apart.

As an example, let us calculate the potential energy of an arrangement of eight negative charges on the corners of a cube of side b, with a positive charge in the center of the cube, as in Fig. 1.6*a*. Suppose each negative charge is an electron with charge $-e$, while the central particle carries a double positive charge, $2e$. Summing over all pairs, we have

$$U = \frac{8(-2e^2)}{(\sqrt{3}/2)b} + \frac{12e^2}{b} + \frac{12e^2}{\sqrt{2}b} + \frac{4e^2}{\sqrt{3}b} = \frac{4.32e^2}{b} \tag{8}$$

Figure 1.6*b* shows where each term in this sum comes from. The energy is positive, indicating that work had to be done on the system to assemble it. That work could, of course, be recovered if we let the charges move apart, exerting forces on some external body or bodies. Or if the electrons were simply to fly apart from this configuration, the *total kinetic energy* of all the particles would become equal to U. This would be true whether they came apart simultaneously and symmetrically, or were released one at a time in any order. Here we see the power of this simple notion of the total potential energy of the system. Think what the problem would be like if we had to compute the resultant vector force on every particle at every stage of assembly of the configuration! In this example, to be sure, the geometrical symmetry would simplify that task; even so, it would be more complicated than the simple calculation above.

One way of writing the instruction for the sum over pairs is this:

$$U = \frac{1}{2}\sum_{j=1}^{N}\sum_{k\neq j}\frac{q_j q_k}{r_{jk}} \tag{9}$$

The double-sum notation, $\Sigma_{j=1}^{N}\Sigma_{k\neq j}$, says: Take $j = 1$ and sum over $k = 2, 3, 4, \ldots, N$; then take $j = 2$ and sum over $k = 1, 3, 4, \ldots, N$; and so on, through $j = N$. Clearly this includes every pair *twice,* and to correct for that we put in front the factor ½.

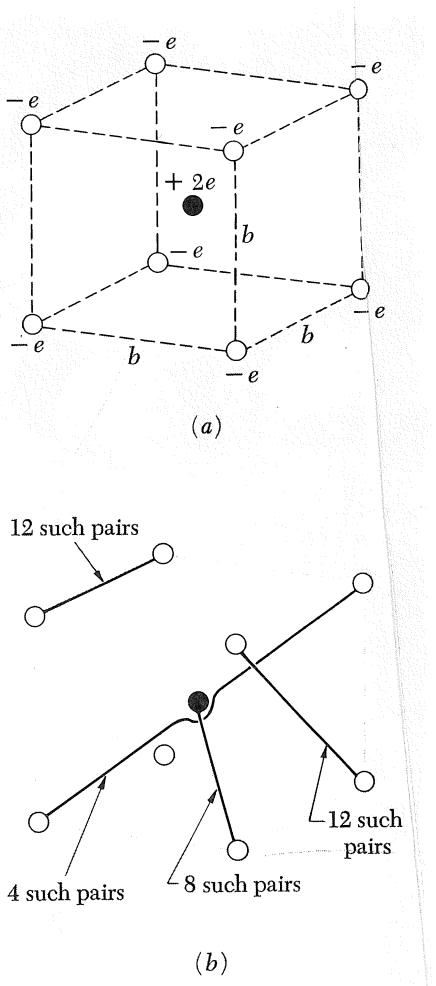

FIGURE 1.6
(*a*) The potential energy of this arrangement of nine point charges is given by Eq. 9. (*b*) Four types of pairs are involved in the sum.

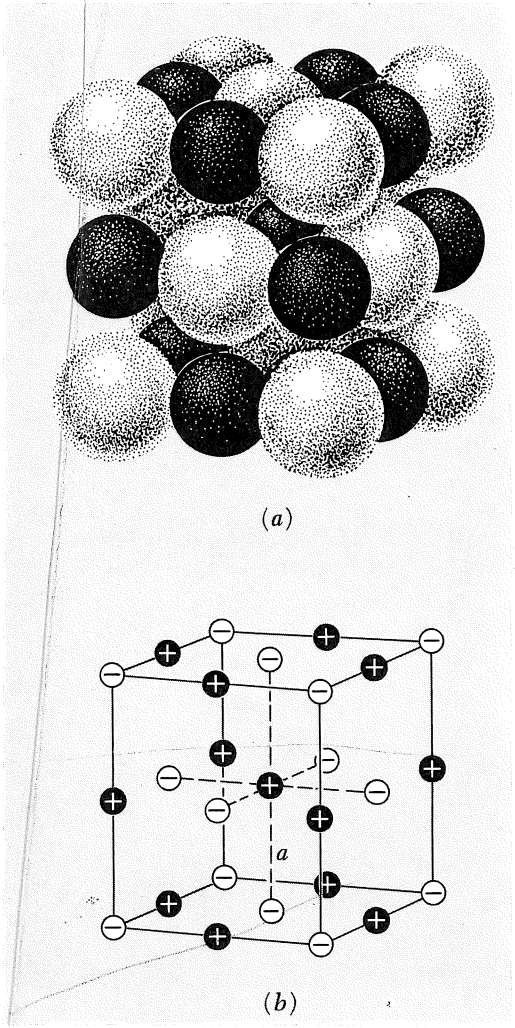

FIGURE 1.7
A portion of a sodium chloride crystal, with the ions Na^+ and $C1^-$ shown in about the right relative proportions (*a*), and replaced by equivalent point charges (*b*).

ELECTRICAL ENERGY IN A CRYSTAL LATTICE

1.6 These ideas have an important application in the physics of crystals. We know that an ionic crystal like sodium chloride can be described, to a very good approximation, as an arrangement of positive ions (Na^+) and negative ions (Cl^-) alternating in a regular three-dimensional array or lattice. In sodium chloride the arrangement is that shown in Fig. 1.7*a*. Of course the ions are not point charges, but they are nearly spherical distributions of charge and therefore (as we shall presently prove) the electrical forces they exert on one another are the same as if each ion were replaced by an equivalent point charge at its center. We show this electrically equivalent system in Fig. 1.7*b*. The electrostatic potential energy of the lattice of charges plays an important role in the explanation of the stability and cohesion of the ionic crystal. Let us see if we can estimate its magnitude.

We seem to be faced at once with a sum that is enormous, if not doubly infinite, for any macroscopic crystal contains 10^{20} atoms at least. Will the sum converge? Now what we hope to find is the potential energy per unit volume or mass of crystal. We confidently expect this to be independent of the size of the crystal, based on the general argument that one end of a macroscopic crystal can have little influence on the other. Two grams of sodium chloride ought to have twice the potential energy of 1 gm, and the shape should not be important so long as the surface atoms are a small fraction of the total number of atoms. We would be *wrong* in this expectation if the crystal were made out of ions of one sign only. Then, 1 gm of crystal would carry an enormous electric charge, and putting two such crystals together to make a 2-gm crystal would take a fantastic amount of energy. (You might estimate how much!) The situation is saved by the fact that the crystal structure is an alternation of equal and opposite charges, so that any macroscopic bit of crystal is very nearly neutral.

To evaluate the potential energy we first observe that every positive ion is in a position equivalent to that of every other positive ion. Furthermore, although it is perhaps not immediately obvious from Fig. 1.7, the arrangement of positive ions around a negative ion is exactly the same as the arrangement of negative ions around a positive ion, and so on. Hence we may take one ion as a center, it matters not which kind, sum over *its* interactions with all the others, and simply multiply by the total number of ions of both kinds. This reduces the double sum in Eq. 9, to a single sum and a factor N; we must still apply the factor ½ to compensate for including each pair twice. That is, the energy of a sodium chloride lattice composed of a total of N ions is

$$U = \frac{1}{2} N \sum_{k=2}^{N} \frac{q_1 q_k}{r_{1k}} \tag{10}$$

Taking the positive ion at the center as in Fig. 1.7*b*, our sum runs over all its neighbors near and far. The leading terms start out as follows:

$$U = \frac{1}{2} N \left(-\frac{6e^2}{a} + \frac{12e^2}{\sqrt{2}a} - \frac{8e^2}{\sqrt{3}a} + \cdots \right) \tag{11}$$

The first term comes from the 6 nearest chlorine ions, at distance a, the second from the 12 sodium ions on the cube edges, and so on. It is clear, incidentally, that this series does not converge *absolutely;* if we were so foolish as to try to sum all the positive terms first, that sum would diverge. To evaluate such a sum, we should arrange it so that as we proceed outward, including ever more distant ions, we include them in groups which represent nearly neutral shells of material. Then if the sum is broken off, the more remote ions which have been neglected will be such an even mixture of positive and negative charges that we can be confident their contribution would have been small. This is a crude way to describe what is actually a somewhat more delicate computational problem. The numerical evaluation of such a series is easily accomplished with a computer. The answer in this example happens to be

$$U = \frac{-0.8738 N e^2}{a} \tag{12}$$

Here N, the number of ions, is twice the number of NaCl molecules.

The negative sign shows that work would have to be *done* to take the crystal apart into ions. In other words, the electrical energy helps to explain the cohesion of the crystal. If this were the whole story, however, the crystal would collapse, for the potential energy of the charge distribution is obviously *lowered* by shrinking all the distances. We meet here again the familiar dilemma of classical—that is, non-quantum—physics. No system of stationary particles can be in stable equilibrium, according to classical laws, under the action of electrical forces alone. Does this make our analysis useless? Not at all. Remarkably, and happily, in the quantum physics of crystals the electrical potential energy can still be given meaning, and can be computed very much in the way we have learned here.

THE ELECTRIC FIELD

1.7 Suppose we have some arrangement of charges, $q_1, q_2, \ldots, q_N$, fixed in space, and we are interested not in the forces they exert on one another but only in their effect on some other charge q_0 which might be brought into their vicinity. We know how to calculate the

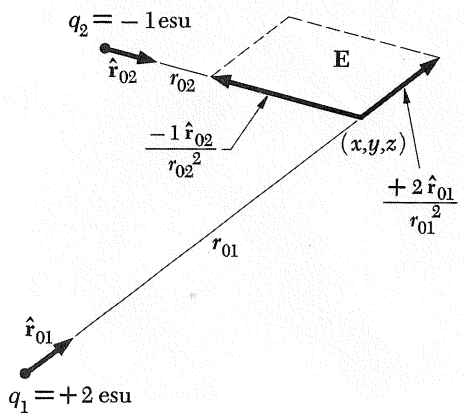

FIGURE 1.8
The field at a point is the vector sum of the fields of each of the charges in the system.

resultant force on this charge, given its position which we may specify by the coordinates x, y, z. The force on the charge q_0 is

$$\mathbf{F}_0 = \sum_{j=1}^{N} \frac{q_0 q_j \hat{\mathbf{r}}_{0j}}{r_{0j}^2} \tag{13}$$

where $\mathbf{r}_{0j}$ is the vector from the jth charge in the system to the point (x, y, z). The force is proportional to q_0, so if we divide out q_0 we obtain a vector quantity which depends only on the structure of our original system of charges, $q_1, \ldots, q_N$, and on the position of the point (x, y, z). We call this vector function of x, y, z the *electric field* arising from the $q_1, \ldots, q_N$ and use the symbol $\mathbf{E}$ for it. The charges $q_1, \ldots, q_N$ we call *sources* of the field. We may take as the *definition* of the electric field $\mathbf{E}$ of a charge distribution, at the point (x, y, z)

$$\mathbf{E}(x, y, z) = \sum_{j=1}^{N} \frac{q_j \hat{\mathbf{r}}_{0j}}{r_{0j}^2} \tag{14}$$

Figure 1.8 illustrates the vector addition of the field of a point charge of 2 esu to the field of a point charge of -1 esu, at a particular point in space. In the CGS system of units, electric field strength is expressed in dynes per unit charge, that is, dynes/esu.

In SI units with the coulomb as the unit of charge and the newton as the unit of force, the electric field strength $\mathbf{E}$ can be expressed in newtons/coulomb, and Eq. 14 would be written like this:

$$\mathbf{E} = \frac{1}{4\pi\epsilon_0} \sum_{j=1}^{N} \frac{q_j \hat{\mathbf{r}}_{0j}}{r_{0j}^2} \tag{14'}$$

each distance r_{0j} being measured in meters.

After the introduction of the electric potential in the next chapter, we shall have another, and completely equivalent, way of expressing the unit of electric field strength; namely, statvolts/cm in the CGS system of units and volts/meter in SI units.

So far we have nothing really new. The electric field is merely another way of describing the system of charges; it does so by giving the force per unit charge, in magnitude and direction, that an exploring charge q_0 would experience at any point. We have to be a little careful with that interpretation. Unless the source charges are really immovable, the introduction of some finite charge q_0 may cause the source charges to shift their positions, so that the field itself, as defined by Eq. 14, is different. That is why we assumed fixed charges to begin our discussion. People sometimes define the field by requiring q_0 to be an "infinitesimal" test charge, letting $\mathbf{E}$ be the limit of $\mathbf{F}/q_0$ as $q_0 \to 0$. Any flavor of rigor this may impart is illusory. Remember that in the real world we have never observed a charge smaller than e! Actually, if we take Eq. 14 as our *definition* of $\mathbf{E}$, without reference to a test charge, no problem arises and the sources need not be fixed.

If the introduction of a new charge causes a shift in the source charges, then it has indeed brought about a change in the electric field, and if we want to predict the force on the new charge, we must use the new electric field in computing it.

Perhaps you still want to ask, what *is* an electric field? Is it something real, or is it merely a name for a factor in an equation which has to be multiplied by something else to give the numerical value of the force we measure in an experiment? Two observations may be useful here. First, since it works, it doesn't make any difference. That is not a frivolous answer, but a serious one. Second, the fact that the electric field vector at a point in space is all we need know to predict the force that will act on *any* charge at that point is by no means trivial. It might have been otherwise! If no experiments had ever been done, we could imagine that, in two different situations in which unit charges experience equal force, test charges of strength 2 units might experience different forces, depending on the nature of the other charges in the system. If that were true, the field description wouldn't work. The electric field attaches to every point in a system a *local property,* in this sense: If we know **E** in some small neighborhood, we know, *without further inquiry,* what will happen to any charges in that neighborhood. We don't need to ask what produced the field.

To visualize an electric field, you need to associate a vector, that is, a magnitude and direction, with every point in space. We shall use various schemes, none of them wholly satisfactory, to depict vector fields in this book.

It is hard to draw in two dimensions a picture of a vector func-

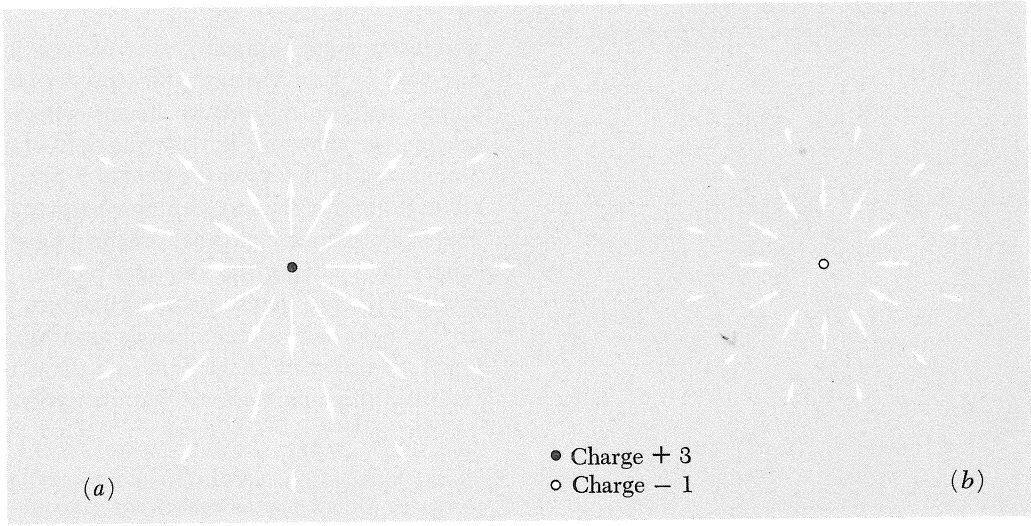

FIGURE 1.9
(*a*) Field of a charge $q_1 = 3$. (*b*) Field of a charge $q_2 = -1$. Both representations are necessarily crude and only roughly quantitative.

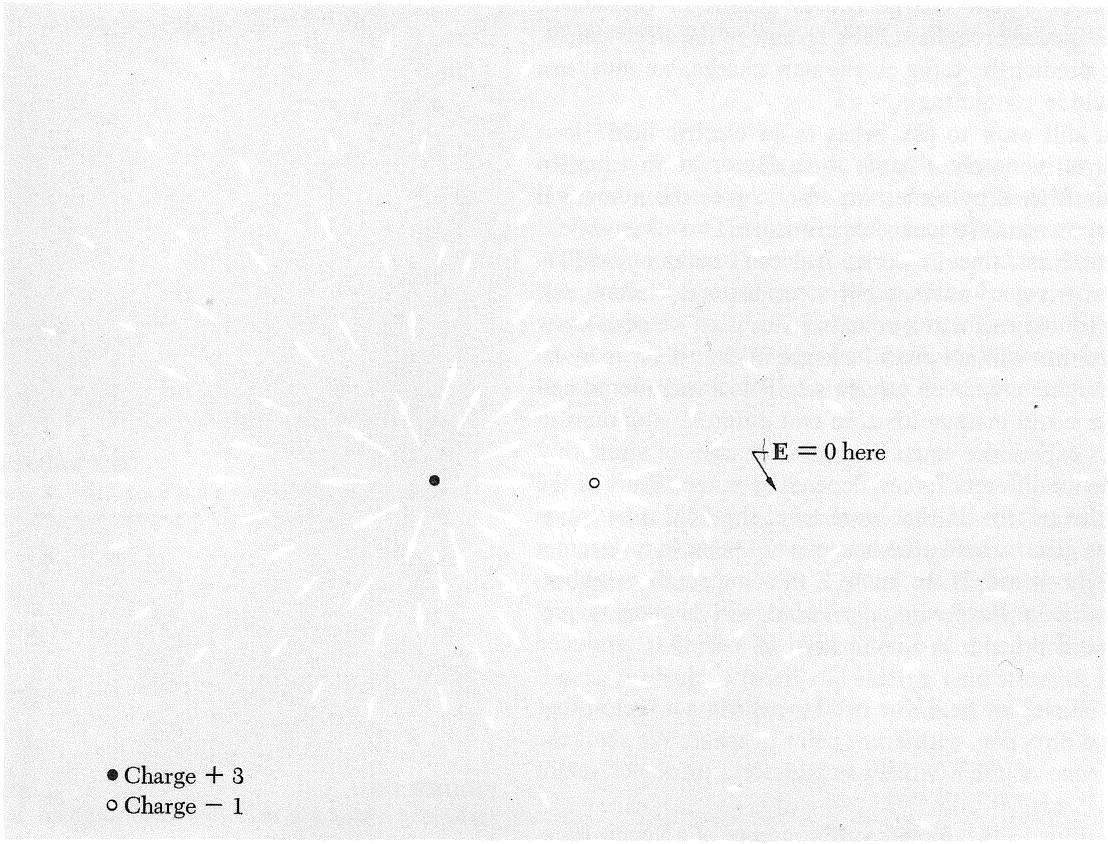

FIGURE 1.10
The field in the vicinity of two charges, $q_1 = +3$, $q_2 = -1$, is the superposition of the fields in Fig. 1.9*a* and *b*.

tion in three-dimensional space. We can indicate the magnitude and direction of **E** at various points by drawing little arrows near those points, making the arrows longer where E is larger.† Using this scheme, we show in Fig. 1.9*a* the field of an isolated point charge of 3 units and in Fig. 1.9*b* the field of a point charge of -1 unit. These pictures admittedly add nothing whatever to our understanding of the field of an isolated charge; anyone can imagine a simple radial inverse-square field without the help of a picture. We show them in order to combine the two fields in Fig. 1.10, which indicates in the same manner the field of two such charges separated by a distance a. All that Fig. 1.10 can show is the field in a plane containing the charges. To get a full three-dimensional representation one must imagine the fig-

†Such a representation is rather clumsy at best. It is hard to indicate the point in space to which a particular vector applies, and the range of magnitudes of E is usually so large that it is impracticable to make the lengths of the arrows proportional to E.

ure rotated around the symmetry axis. In Fig. 1.10 there is one point in space where **E** is zero. How far from the nearest charge must this point lie? Notice also that toward the edge of the picture the field points more or less radially outward all around. One can see that at a very large distance from the charges the field will look very much like the field from a positive point charge. This is to be expected because the separation of the charges cannot make very much difference for points far away, and a point charge of 2 units is just what we would have left if we superimposed our two sources at one spot.

Another way to depict a vector field is to draw *field lines*. These are simply curves whose tangent, at any point, lies in the direction of the field at that point. Such curves will be smooth and continuous

FIGURE 1.11
Some field lines in the electric field around two charges, $q_1 = +3$, $q_2 = -1$.

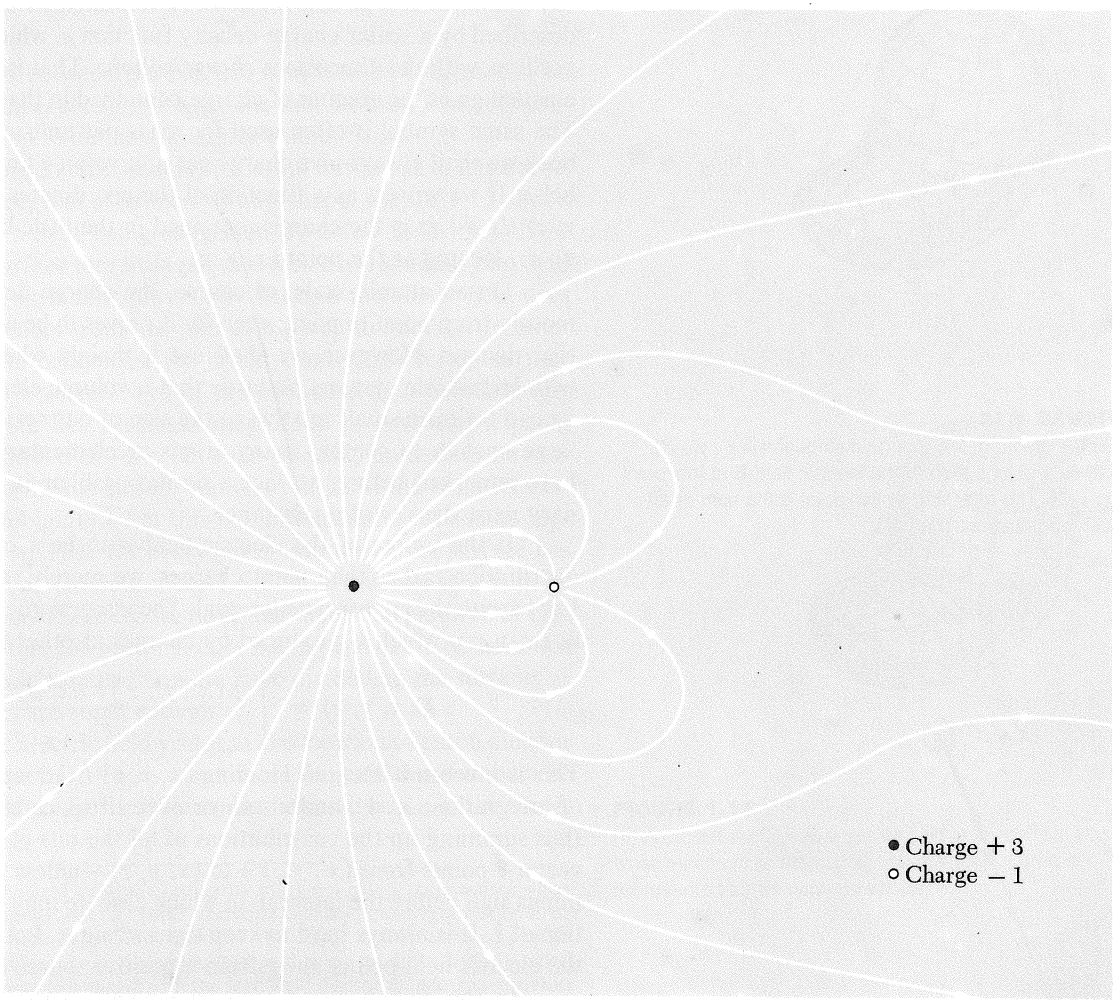

except at singularities such as point charges, or points like the one in the example of Fig. 1.10 where the field is zero. A field line plot does not directly give the magnitude of the field, although we shall see that, in a general way, the field lines converge as we approach a region of strong field and spread apart as we approach a region of weak field. In Fig. 1.11 are drawn some field lines for the same arrangement of charges as in Fig. 1.10, a positive charge of 3 units and a negative charge of 1 unit. Again, we are restricted by the nature of paper and ink to a two-dimensional section through a three-dimensional bundle of curves.

CHARGE DISTRIBUTIONS

1.8 This is as good a place as any to generalize from *point charges* to *continuous charge distributions*. A volume distribution of charge is described by a scalar charge-density function ρ, which is a function of position, with the dimensions *charge/volume*. That is, ρ times a volume element gives the amount of charge contained in that volume element. The same symbol is often used for mass per unit volume, but in this book we shall always give charge per unit volume first call on the symbol ρ. If we write ρ as a function of the coordinates x, y, z, then $\rho(x, y, z)\ dx\ dy\ dz$ is the charge contained in the little box, of volume $dx\ dy\ dz$, located at the point (x, y, z).

On an atomic scale, of course, the charge density varies enormously from point to point; even so, it proves to be a useful concept in that domain. However, we shall use it mainly when we are dealing with large-scale systems, so large that a volume element $dv = dx\ dy\ dz$ can be quite small relative to the size of our system, although still large enough to contain many atoms or elementary charges. As we have remarked before, we face a similar problem in defining the ordinary mass density of a substance.

If the source of the electric field is to be a continuous charge distribution rather than point charges, we merely replace the sum in Eq. 14 with the appropriate integral. The integral gives the electric field at (x, y, z), which is produced by charges at other points (x', y', z').

$$\mathbf{E}(x, y, z) = \int \frac{\rho(x', y', z')\,\hat{\mathbf{r}}\,dx'\,dy'\,dz'}{r^2} \tag{15}$$

This is a volume integral. Holding (x, y, z) fixed we let the variables of integration, x', y', and z', range over all space containing charge, thus summing up the contributions of all the bits of charge. The unit vector $\hat{\mathbf{r}}$ points from (x', y', z') to (x, y, z)—unless we want to put a minus sign before the integral, in which case we may reverse the direction of $\hat{\mathbf{r}}$. It is always hard to keep signs straight. Let's remember that the electric field points *away* from a positive source (Fig. 1.12).

In the neighborhood of a true point charge the electric field

FIGURE 1.12
Each element of the charge distribution $\rho(x', y', z')$ makes a contribution to the electric field **E** at this point (x, y, z). The total field at this point is the sum of all such contributions (Eq. 15).

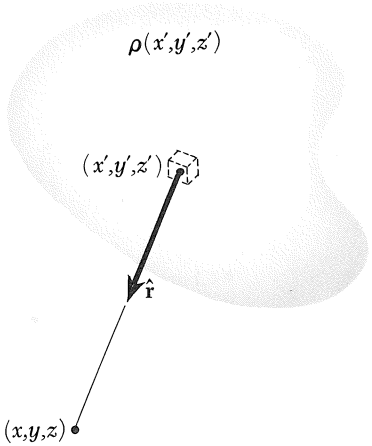

grows infinite like $1/r^2$ as we approach the point. It makes no sense to talk about the field *at* the point charge. As our ultimate physical sources of field are not, we believe, infinite concentrations of charge in zero volume but instead finite structures, we simply ignore the mathematical singularities implied by our point-charge language and rule out of bounds the interior of our elementary sources. A continuous charge distribution $\rho(x', y', z')$ which is nowhere infinite gives no trouble at all. Equation 15 can be used to find the field at any point within the distribution. The integrand doesn't blow up at $r = 0$ because the volume element in the numerator is in that limit proportional to $r^2\,dr$. That is to say, so long as ρ remains finite, the field will remain finite everywhere, even in the interior or on the boundary of a charge distribution.

FLUX

1.9 The relation between the electric field and its sources can be expressed in a remarkably simple way, one that we shall find very useful. For this we need to define a quantity called *flux*.

Consider some electric field in space and in this space some arbitrary closed surface, like a balloon of any shape. Figure 1.13 shows such a surface, the field being suggested by a few field lines. Now divide the whole surface into little patches which are so small that over any one patch the surface is practically flat and the vector field does not change appreciably from one part of a patch to another. In other words, don't let the balloon be too crinkly, and don't let its surface pass right through a singularity† of the field such as a point charge. The area of a patch has a certain magnitude in cm^2, and a patch defines a unique direction—the outward-pointing normal to its surface. (Since the surface is closed, you can tell its inside from its outside; there is no ambiguity.) Let this magnitude and direction be represented by a vector. Then for every patch into which the surface has been divided, such as patch number j, we have a vector $\mathbf{a}_j$ giving its area and orientation. The steps we have just taken are pictured in Fig. 1.13*b* and *c*. Note that the vector $\mathbf{a}_j$ does not depend at all on the shape of the patch; it doesn't matter how we have divided up the surface, as long as the patches are small enough.

Let $\mathbf{E}_j$ be the electric field vector at the location of patch number j. The scalar product $\mathbf{E}_j \cdot \mathbf{a}_j$ is a number. We call this number the *flux* through that bit of surface. To understand the origin of the name,

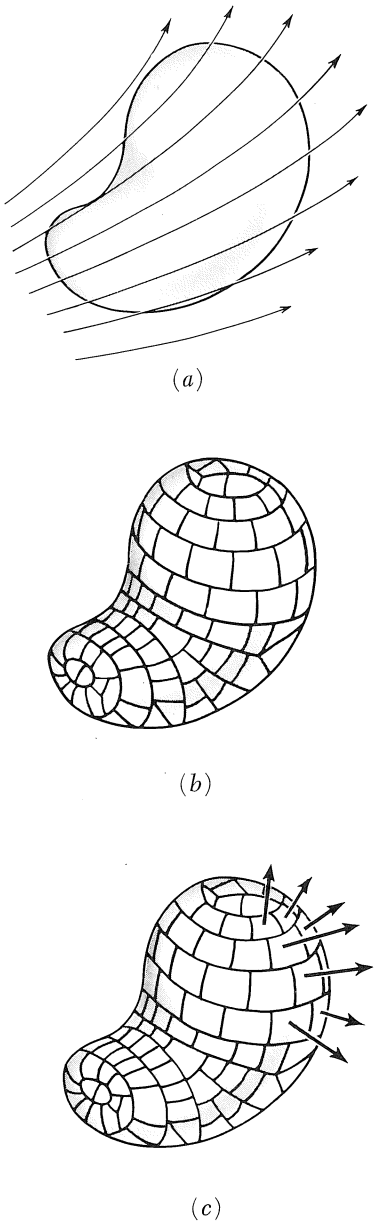

FIGURE 1.13
(*a*) A closed surface in a vector field is divided (*b*) into small elements of area. (*c*) Each element of area is represented by an outward vector.

†By a singularity of the field we would ordinarily mean not only a point source where the field approaches infinity, but any place where the field changes magnitude or direction discontinuously, such as an infinitesimally thin layer of concentrated charge. Actually this latter, milder, kind of singularity would cause no difficulty here unless our balloon's surface were to coincide with the surface of discontinuity over some finite area.